Contact homology and virtual fundamental cycles

John Pardon*

28 October 2015

Abstract

We give a construction of contact homology in the sense of Eliashberg–Givental–Hofer. Specifically, we construct coherent virtual fundamental cycles on the relevant compactified moduli spaces of pseudo-holomorphic curves.

The aim of this work is to provide a rigorous construction of *contact homology*, an invariant of contact manifolds and exact symplectic cobordisms due to Eliashberg–Givental–Hofer [Eli98, EGH00]. The contact homology of (Y, ξ) is defined in terms of pseudo-holomorphic curve "counts" (in the sense of Gromov [Gro85]) in the symplectization $\mathbb{R} \times Y$. Unfortunately, the moduli spaces of such curves often suffer from a severe lack of transversality, and for this reason, it has, until now, remained an important open problem to define the relevant curve counts in full generality. We will use the framework developed in [Par15] to construct coherent virtual fundamental cycles on the moduli spaces in question, thus giving rise to the desired curve counts. Our methods are quite general, and should apply equally well to many other moduli spaces of interest.

Previously, a number of important special cases of contact homology (and closely related invariants) have been constructed rigorously using generic and/or automatic transversality techniques. Cylindrical contact homology of some three-manifolds was constructed by Bao–Honda [BH15] and Hutchings–Nelson [HN15]. Legendrian contact homology in $\mathbb{R}^{2n+1}$ was constructed by Ekholm–Etnyre–Sullivan [EES05]. Embedded contact homology was introduced and constructed by Hutchings and Hutchings–Taubes [Hut02, Hut09, HT07, HT09a].

Remark 0.1 (Virtual moduli cycle techniques). The technique of patching together local finite-dimensional reductions to construct virtual fundamental cycles has been applied to moduli spaces of pseudo-holomorphic curves by many authors, including Fukaya–Ono [FO99] (Kuranishi structures), Li–Tian [LT98a], Liu–Tian [LT98b], Ruan [Rua99], Fukaya–Oh–Ohta–Ono [FO99, FOOO09a, FOOO09b, FOOO12, FOOO15], McDuff–Wehrheim [MW15c, MW15a, MW15b] (Kuranishi atlases), Joyce [Joy15b, Joy14, Joy12] (Kuranishi spaces and d-manifolds), and [Par15] (implicit atlases).

More recently, the theory of polyfolds developed by Hofer–Wysocki–Zehnder [HWZ07, HWZ09a, HWZ09b, HWZ10a, HWZ10b, HWZ11, HWZ14] appears to provide a robust new

*This research was partially conducted during the period the author served as a Clay Research Fellow. The author was also partially supported by a National Science Foundation Graduate Research Fellowship under grant number DGE–1147470.

infinite-dimensional context in which all reasonable moduli spaces of pseudo-holomorphic curves may be perturbed "abstractly" to obtain transversality.

Any one of the above theories (once sufficiently developed) could be used to prove the main results of this paper (we use the theory from [Par15]). Although these theories vary in their approach to the myriad of technical issues involved, they are expected to give rise to completely equivalent virtual fundamental cycles.

Remark 0.2 (Historical discussion). The theory of pseudo-holomorphic curves in closed symplectic manifolds was founded by Gromov [Gro85]. Hofer's breakthrough work on the three-dimensional Weinstein conjecture [Hof93] introduced pseudo-holomorphic curves in symplectizations and their relation with Reeb dynamics. The analytic theory of such curves was then further developed by Hofer–Wysocki–Zehnder [HWZ96, HWZ98a, HWZ95, HWZ99, HWZ02]. On the algebraic side, Eliashberg–Givental–Hofer [EGH00] introduced the theories of contact homology and symplectic field theory, based on counts of pseudo-holomorphic curves in symplectizations and symplectic manifolds with symplectization ends (assuming such counts can be defined). The key compactness results for moduli spaces of such curves were established by Bourgeois–Eliashberg–Hofer–Wyzocki–Zehnder [BEHWZ03]. Gluing techniques applicable to such pseudo-holomorphic curves have been developed by many authors, notably Taubes, Donaldson, Floer, Fukaya–Oh–Ohta–Ono, and Hofer–Wysocki–Zehnder.

Acknowledgements: The author thanks Yasha Eliashberg for introducing the author to this problem and for many useful conversations. The author is also grateful for useful discussions with Mohammed Abouzaid, Frédéric Bourgeois, Roger Casals, Vincent Colin, Tobias Ekholm, Kenji Fukaya, Eleny Ionel, Patrick Massot, Rafe Mazzeo, Peter Ozsváth, and Chris Wendl. A preliminary version of this paper appeared as part of the author's 2015 Ph.D. thesis at Stanford University supervised by Ya. Eliashberg.

1 Statement of results

We begin by stating our main results while simultaneously reviewing the definition of contact homology (for additional context, the reader may consult [Eli98, EGH00, Eli07]). Specifically, we discuss the four components of the definition: (I) the differential, (II) the cobordism map, (III) the deformation chain homotopy, and (IV) the composition chain homotopy. We then discuss how to assemble these components into the contact homology functor. We also review some applications, and we provide some technical remarks for the experts.

In the later sections of the paper, we invite the reader to restrict their attention to (I) on a first reading. The generalization from (I) to (II), (III), (IV) is mostly straightforward, and may be saved for a second reading.

1.1 (I) The differential

We will show how to define the contact homology differential from the following setup.

Setup I. This setup consists of a triple (Y, λ, J) as follows. Here Y^{2n-1} is a closed manifold[1], and λ is a *contact form* on Y (i.e. a 1-form such that $\lambda \wedge (d\lambda)^{n-1}$ is nonvanishing). We denote by $\xi := \ker \lambda$ the induced co-oriented contact structure, and $J : \xi \to \xi$ is an almost complex structure which is *compatible* with $d\lambda$ (i.e. $d\lambda(\cdot, J\cdot)$ is a positive definite symmetric pairing on ξ).

Denote by R_λ the *Reeb vector field* associated to λ (defined by the properties $\lambda(R_\lambda) = 1$ and $d\lambda(R_\lambda, \cdot) = 0$). We will denote by $\mathcal{P} = \mathcal{P}(Y, \lambda)$ the collection of (unparameterized) *Reeb orbits* (i.e. closed trajectories of R_λ, not necessarily embedded). We make the additional assumption in this setup that all Reeb orbits are *non-degenerate* (i.e. the linearized return map has no fixed vector).

Let $\hat{Y} := \mathbb{R} \times Y$ (with coordinate $s \in \mathbb{R}$) denote the *symplectization*[2] of Y. Now J induces an $\mathbb{R}$-invariant almost complex structure $\hat{J}$ on $\hat{Y}$ defined by the property that $\hat{J}(\partial_s) = R_\lambda$ and $\hat{J}|_\xi = J$. Given a Reeb orbit $\gamma^+ \in \mathcal{P}$ and a finite set of Reeb orbits $\Gamma^- \to \mathcal{P}$, we define:

$$\pi_2(Y, \gamma^+ \sqcup \Gamma^-) := [(S, \partial S), (Y, \gamma^+ \sqcup \Gamma^-)]/\mathrm{Aut}(S, \partial S) \tag{1.1}$$

where S is any compact connected oriented surface of genus zero with boundary, equipped with a homeomorphism between ∂S and $\gamma^+ \sqcup \Gamma^-$ (preserving orientation on γ^+ and reversing orientation on Γ^-). There is a natural partition $\mathcal{P} = \mathcal{P}_{\mathrm{good}} \sqcup \mathcal{P}_{\mathrm{bad}}$, and for each good Reeb orbit $\gamma \in \mathcal{P}_{\mathrm{good}}$, there is an associated *orientation line* (i.e. a $\mathbb{Z}/2$-graded free $\mathbb{Z}$-module of rank one) $\mathfrak{o}_\gamma$ with parity $|\gamma| := \mathrm{sign}(\det(I - A_\gamma)) \in \{\pm 1\} = \mathbb{Z}/2$, where A_γ denotes the linearized return map of γ acting on ξ (see §2.5). We set[3] $\mathfrak{o}_\Gamma := \bigotimes_{\gamma \in \Gamma} \mathfrak{o}_\gamma$ and $|\Gamma| := \sum_{\gamma \in \Gamma} |\gamma|$ for any finite set of Reeb orbits $\Gamma \to \mathcal{P}$. For a given Reeb orbit $\gamma \in \mathcal{P}$, let $d_\gamma \in \mathbb{Z}_{\geq 1}$ denote its covering multiplicity.

Let $\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^-, \Gamma^+; \beta)$ denote the compactified moduli space of connected $\hat{J}$-holomorphic curves of genus zero in $\hat{Y}$ modulo $\mathbb{R}$-translation, with one positive puncture asymptotic to γ^+ and negative punctures asymptotic to Γ^-, in the homotopy class β, along with asymptotic markers on the domain mapping to fixed basepoints on γ^+ and Γ^- (see §§2.1–2.2). We denote by $\mu(\gamma^+, \Gamma^-; \beta) \in \mathbb{Z}$ the *index* of this moduli problem (the "virtual" or "expected" dimension of $\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+, \Gamma^-; \beta)$ equals $\mu(\gamma^+, \Gamma^-; \beta) - 1$); we have $\mu(\gamma^+, \Gamma^-; \beta) \equiv |\gamma^+| - |\Gamma^-| \in \mathbb{Z}/2$ (see §2.4). We say that $\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+, \Gamma^-; \beta)$ is *regular* (or *cut out transversally*) iff the relevant linearized operator is everywhere surjective (see §2.3). By [BEHWZ03], each $\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+, \Gamma^-; \beta)$ is compact.

Our main result in this setup is the following.

Theorem I. *Fix (Y, λ, J) as in Setup I. There exists a non-empty set $\Theta_{\mathrm{I}} = \Theta_{\mathrm{I}}(Y, \lambda, J)$ along with, for all $\theta \in \Theta_{\mathrm{I}}$, $\gamma^+ \in \mathcal{P}_{\mathrm{good}}$, $\Gamma^- \to \mathcal{P}_{\mathrm{good}}$, and $\beta \in \pi_2(Y, \gamma^+ \sqcup \Gamma^-)$, numbers ("virtual*

[1] Everything is in the smooth category unless stated otherwise.

[2] More intrinsically, the symplectization of a co-oriented contact manifold (Y, ξ) is defined as the total space of the bundle of 1-forms with kernel ξ, namely $\hat{Y} := \ker(T^*Y \to \xi^*)_+$. The restriction of the tautological Liouville 1-form on T^*Y is a Liouville 1-form $\hat{\lambda}$ on $\hat{Y}$; the associated Liouville vector field on $\hat{Y}$ generates an $\mathbb{R}$-action on $\hat{Y}$ which is simply scaling by e^s. A choice of contact form λ for ξ induces an identification of $(\hat{Y}, \hat{\lambda})$ with $(\mathbb{R} \times Y, e^s\lambda)$.

[3] By $\otimes$ we always mean the super tensor product, i.e. where the isomorphism $A \otimes B \xrightarrow{\sim} B \otimes A$ is given by $a \otimes b \mapsto (-1)^{|a||b|} b \otimes a$.

moduli counts"):[4]

$$\#\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)^{\mathrm{vir}}_{\theta} \in \mathfrak{o}^\vee_{\gamma^+}\otimes\mathfrak{o}_{\Gamma^-}\otimes\mathbb{Q} \tag{1.2}$$

This data is functorial[5] *in (Y,λ,J), and the following properties are satisfied for all $\theta\in\Theta_{\mathrm{I}}$:*

i. If $\mu(\gamma^+,\Gamma^-;\beta)\neq 1$, then $\#\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)^{\mathrm{vir}}_{\theta}=0$.

ii. If $\mu(\gamma^+,\Gamma^-;\beta)=1$ and $\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)$ is regular, then $\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)=\mathcal{M}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)$ is a compact manifold of dimension zero naturally oriented by $\mathfrak{o}_{\gamma^+}\otimes\mathfrak{o}^\vee_{\Gamma^-}$, and:

$$\#\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)^{\mathrm{vir}}_{\theta}=\#\mathcal{M}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta) \tag{1.3}$$

(in particular, if $\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)=\varnothing$ then $\#\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)^{\mathrm{vir}}_{\theta}=0$).

iii. The virtual moduli counts satisfy the following "master equation":

$$\#\partial\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)^{\mathrm{vir}}_{\theta}=0 \tag{1.4}$$

for all $(\gamma^+,\Gamma^-;\beta)$, where the left hand side denotes the sum over all codimension one boundary strata of the relevant products of (1.2) and inverse covering multiplicities of intermediate orbits (this sum is finite by compactness).

For (Y,λ) as in Setup I, let:

$$CC_\bullet(Y,\xi)_\lambda := \bigoplus_{n\geq 0}\mathrm{Sym}^n_{\mathbb{Q}}\Big(\bigoplus_{\gamma\in\mathcal{P}_{\mathrm{good}}}\mathfrak{o}_\gamma\Big) \tag{1.5}$$

denote the free supercommutative (i.e. $ab=(-1)^{|a||b|}ba$) unital $\mathbb{Z}/2$-graded $\mathbb{Q}$-algebra generated by $\mathfrak{o}_\gamma$ for $\gamma\in\mathcal{P}_{\mathrm{good}}$. Given J and $\theta\in\Theta_{\mathrm{I}}(Y,\lambda,J)$, we may define a differential:

$$d_{J,\theta}: CC_\bullet(Y,\xi)_\lambda\to CC_{\bullet-1}(Y,\xi)_\lambda \tag{1.6}$$

which satisfies the Leibniz rule ($d(1)=0$ and $d(ab)=da\cdot b+(-1)^{|a|}a\cdot db$); it is defined by the property that it acts on $\mathfrak{o}_{\gamma^+}$ by pairing on the left with $d^{-1}_{\gamma^+}$ times the sum over all (Γ^-,β) of (1.2) divided by $\#|\mathrm{Aut}(\Gamma^-,\beta)|$. By [BEHWZ03], for fixed γ^+ there are only finitely many non-empty spaces $\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)$, so $d_{J,\theta}$ is well-defined. The master equation (1.4) implies that this differential squares to zero:

$$d_{J,\theta}d_{J,\theta}=0 \tag{1.7}$$

We denote the resulting homology by:

$$CH_\bullet(Y,\xi)_{\lambda,J,\theta} \tag{1.8}$$

which is itself a supercommutative unital $\mathbb{Z}/2$-graded $\mathbb{Q}$-algebra.

Remark 1.1. One can view contact homology as a version of S^1-equivariant Morse–Floer homology of the loop space LY with respect to the action functional $\gamma\mapsto\int_\gamma\lambda$.

[4] Note that for $\mathfrak{o}$ odd, there is no symmetric perfect pairing $\mathfrak{o}\otimes\mathfrak{o}\to\mathbb{Z}$, so we may not conflate $\mathfrak{o}$ and its dual $\mathfrak{o}^\vee$.

[5] Functoriality has the usual meaning here: an isomorphism between triples $i:(Y,\lambda,J)\xrightarrow{\sim}(Y',\lambda',J')$ induces an isomorphism $i_*:\Theta_{\mathrm{I}}(Y,\lambda,J)\xrightarrow{\sim}\Theta_{\mathrm{I}}(Y',\lambda',J')$ such that $\mathrm{id}_*=\mathrm{id}$ and $(i\circ j)_*=i_*\circ j_*$, and $\#\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+,\Gamma^-;\beta)^{\mathrm{vir}}_{\theta}=\#\overline{\mathcal{M}}_{\mathrm{I}}(i_*\gamma^+,i_*\Gamma^-;i_*\beta)^{\mathrm{vir}}_{i_*\theta}$.

1.2 (II) The cobordism map

We will show how to define the contact homology cobordism map from the following setup.

Setup II. This setup consists of a triple $(\hat{X}, \hat{\omega}, \hat{J})$ and two triples $(Y^\pm, \lambda^\pm, J^\pm)$ as in Setup I. Here $\hat{X}^{2n}$ is a manifold and $\hat{\omega}$ is a symplectic form on $\hat{X}$ (i.e. a 2-form such that $d\hat{\omega} = 0$ and $\hat{\omega}^n$ is non-vanishing). Here $\hat{J}$ is an almost complex structure on $\hat{X}$ which is *tamed* by $\hat{\omega}$, that is $\hat{\omega}(v, \hat{J}v) > 0$ for nonzero $v \in T\hat{X}$. This setup also includes the data of proper maps:

$$([0,\infty) \times Y^+, d(e^s\lambda^+), \hat{J}^+) \to (\hat{X}, \hat{\omega}, \hat{J}) \tag{1.9}$$

$$((-\infty,0] \times Y^-, d(e^s\lambda^-), \hat{J}^-) \to (\hat{X}, \hat{\omega}, \hat{J}) \tag{1.10}$$

which are diffeomorphisms onto their (not necessarily disjoint) images, and we require that their image in $\hat{X}$ have precompact complement. The charts (1.9)–(1.10) need only be defined for $|s|$ sufficiently large, and they need only preserve symplectic forms up to a constant scaling factor.

Let $\mathcal{P}^\pm := \mathcal{P}(Y^\pm, \lambda^\pm)$. Given a Reeb orbit $\gamma^+ \in \mathcal{P}^+$, a finite set of Reeb orbits $\Gamma^- \to \mathcal{P}^-$, and a homotopy class $\beta \in \pi_2(\hat{X}, \gamma^+ \sqcup \Gamma^-)$, we denote by $\overline{\mathcal{M}}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta)$ the compactified moduli space of connected $\hat{J}$-holomorphic curves of genus zero in $\hat{X}$ from γ^+ to Γ^- in the homotopy class β. See §§2.1–2.2 for more details. By [BEHWZ03], each $\overline{\mathcal{M}}_{\mathrm{I}}(\gamma^+, \Gamma^-; \beta)$ is compact.

Theorem II. *Fix data as in Setup II, and let $\Theta_{\mathrm{I}}^\pm = \Theta_{\mathrm{I}}(Y^\pm, \lambda^\pm, J^\pm)$. There exists a set Θ_{II} with a surjective map $\Theta_{\mathrm{II}} \twoheadrightarrow \Theta_{\mathrm{I}}^+ \times \Theta_{\mathrm{I}}^-$, along with, for all $\theta \in \Theta_{\mathrm{II}}$, $\gamma^+ \in \mathcal{P}^+_{\mathrm{good}}$, $\Gamma^- \to \mathcal{P}^-_{\mathrm{good}}$, and $\beta \in \pi_2(X, \gamma^+ \sqcup \Gamma^-)$, numbers ("virtual moduli counts"):*

$$\#\overline{\mathcal{M}}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta)^{\mathrm{vir}}_\theta \in \mathfrak{o}^\vee_{\gamma^+} \otimes \mathfrak{o}_{\Gamma^-} \otimes \mathbb{Q} \tag{1.11}$$

This data is functorial in the data from Setup II, and the following properties are satisfied for all $\theta \in \Theta_{\mathrm{II}}$:

i. If $\mu(\gamma^+, \Gamma^-; \beta) \neq 0$, then $\#\overline{\mathcal{M}}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta)^{\mathrm{vir}}_\theta = 0$.

ii. If $\mu(\gamma^+, \Gamma^-; \beta) = 0$ and $\overline{\mathcal{M}}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta)$ is regular, then $\overline{\mathcal{M}}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta) = \mathcal{M}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta)$ is a compact manifold of dimension zero naturally oriented by $\mathfrak{o}_{\gamma^+} \otimes \mathfrak{o}^\vee_{\Gamma^-}$, and:

$$\#\overline{\mathcal{M}}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta)^{\mathrm{vir}}_\theta = \#\mathcal{M}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta) \tag{1.12}$$

iii. The numbers $\#\overline{\mathcal{M}}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta)^{\mathrm{vir}}_\theta$ satisfy the following "master equation":

$$\#\partial\overline{\mathcal{M}}_{\mathrm{II}}(\gamma^+, \Gamma^-; \beta)^{\mathrm{vir}}_\theta = 0 \tag{1.13}$$

for all $(\gamma^+, \Gamma^-; \beta)$, where the left hand side denotes the sum over all codimension one boundary strata of the relevant products of (1.2) (using $\theta^\pm$), (1.11), and inverse covering multiplicities of intermediate orbits (this sum is finite by compactness).

Let $(\hat{X}, d\hat{\lambda}, \hat{J})$ and $(Y^\pm, \lambda^\pm, J^\pm)$ be as in Setup II, and let $\theta \in \Theta_{\mathrm{II}}$. We may define a unital $\mathbb{Q}$-algebra map:

$$\Phi(\hat{X}, \hat{\lambda})_{\hat{J},\theta} : CC_\bullet(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \to CC_\bullet(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-} \tag{1.14}$$

by the property that it acts on $\mathfrak{o}_{\gamma^+}$ by pairing on the left with $d_{\gamma^+}^{-1}$ times the sum over all (Γ^-,β) of (1.11) divided by $\#|\operatorname{Aut}(\Gamma^-,\beta)|$. By [BEHWZ03] and exactness of $(\hat{X},d\hat{\lambda})$, for fixed γ^+ there are only finitely many non-empty spaces $\overline{\mathcal{M}}_{\mathrm{II}}(\gamma^+,\Gamma^-;\beta)$, so $\Phi(\hat{X},\hat{\lambda})_{\hat{J},\theta}$ is well-defined. Now the master equation (1.13) implies that this is a chain map:

$$d_{J^-,\theta^-}\Phi(\hat{X},\hat{\lambda})_{\hat{J},\theta}-\Phi(\hat{X},\hat{\lambda})_{\hat{J},\theta}d_{J^+,\theta^+}=0 \tag{1.15}$$

We denote the resulting unital $\mathbb{Q}$-algebra map on homology by:

$$\Phi(\hat{X},\hat{\lambda})_{\hat{J},\theta}:CH_\bullet(Y^+,\xi^+)_{\lambda^+,J^+,\theta^+}\to CH_\bullet(Y^-,\xi^-)_{\lambda^-,J^-,\theta^-} \tag{1.16}$$

1.3 (III) The deformation homotopy

Setup III. This setup consists of a triple $(\hat{X},\hat{\omega}^t,\hat{J}^t)_{t\in[0,1]}$ and two triples $(Y^\pm,\lambda^\pm,J^\pm)$, which is a continuous deformation of data as in Setup II. This deformation must be fixed outside a compact subset of $\hat{X}$.

By considering appropriate moduli spaces $\overline{\mathcal{M}}_{\mathrm{III}}(\{\gamma_i^+,\Gamma_i^-;\beta_i\}_{i\in I})$ of (possibly disconnected) $\hat{J}^t$-holomorphic curves in $\hat{X}$, we prove the following result.

Theorem III. *Fix data as in Setup III; let $\Theta_{\mathrm{I}}^\pm=\Theta_{\mathrm{I}}(Y^\pm,\lambda^\pm,J^\pm)$ and let $\Theta_{\mathrm{II}}^{0,1}=\Theta_{\mathrm{II}}(\hat{X},\hat{\lambda}^{0,1},\hat{J}^{0,1})$. There exists a set Θ_{III} along with a surjective map $\Theta_{\mathrm{III}}\twoheadrightarrow\Theta_{\mathrm{II}}^0\times_{\Theta_{\mathrm{I}}^+\times\Theta_{\mathrm{I}}^-}\Theta_{\mathrm{II}}^1$ and virtual moduli counts:*

$$\#\overline{\mathcal{M}}_{\mathrm{III}}(\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I})_\theta^{\mathrm{vir}}\in\left(\bigotimes_{i\in I}\mathfrak{o}_{\gamma_i^+}^\vee\otimes\mathfrak{o}_{\Gamma_i^-}\right)\otimes\mathbb{Q} \tag{1.17}$$

This data is functorial in the data from Setup III, and the following properties are satisfied for all $\theta\in\Theta_{\mathrm{III}}$:

i. If $\sum_{i\in I}\mu(\gamma_i^+,\Gamma_i^-,\beta_i)\neq-1$, then $\#\overline{\mathcal{M}}_{\mathrm{III}}(\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I})_\theta^{\mathrm{vir}}=0$.

ii. If $\sum_{i\in I}\mu(\gamma_i^+,\Gamma_i^-,\beta_i)=-1$ and $\overline{\mathcal{M}}_{\mathrm{III}}(\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I})$ is regular, then $\overline{\mathcal{M}}_{\mathrm{III}}(\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I})=\mathcal{M}_{\mathrm{III}}(\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I})$ is a compact manifold of dimension zero naturally oriented by $\bigotimes_{i\in I}\mathfrak{o}_{\gamma_i^+}\otimes\mathfrak{o}_{\Gamma_i^-}^\vee$, and:

$$\#\overline{\mathcal{M}}_{\mathrm{III}}(\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I})_\theta^{\mathrm{vir}}=\#\mathcal{M}_{\mathrm{III}}(\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I}) \tag{1.18}$$

iii. The numbers $\#\overline{\mathcal{M}}_{\mathrm{III}}(\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I})_\theta^{\mathrm{vir}}$ satisfy the following "master equation":

$$\#\partial\overline{\mathcal{M}}_{\mathrm{III}}(\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I})_\theta^{\mathrm{vir}}=0 \tag{1.19}$$

for all $\{(\gamma_i^+,\Gamma_i^-,\beta_i)\}_{i\in I}$, where the left hand side denotes the sum over all codimension one boundary strata of the relevant products of (1.2) (using $\theta^\pm$), (1.11) (using $\theta^{0,1}$), (1.17), and inverse covering multiplicities of intermediate orbits (this sum is finite by compactness).

For data as in Setup III, it follows that the following two maps coincide:

$$CH_\bullet(Y^+,\xi^+)_{\lambda^+,J^+,\theta^+}\underset{\Phi(\hat{X},\hat{\lambda}^1)_{\hat{J}^1,d\theta^1}}{\overset{\Phi(\hat{X},d\hat{\lambda}^0)_{\hat{J}^0,\theta^0}}{\rightrightarrows}}CH_\bullet(Y^-,\xi^-)_{\lambda^-,J^-,\theta^-} \tag{1.20}$$

for $(\theta^\pm, \theta^{0,1}) \in \Theta^0_{\mathrm{II}} \times_{\Theta^+_{\mathrm{I}} \times \Theta^-_{\mathrm{I}}} \Theta^1_{\mathrm{II}}$. More precisely, one obtains a chain homotopy $K(\hat{X}, \{d\hat{\lambda}^t\}_t)_{\{\hat{J}^t\}_t,\theta} : CC_\bullet(Y^+, \xi^+)_{\lambda^+,J^+,\theta^+} \to CC_\bullet(Y^-, \xi^-)_{\lambda^-,J^-,\theta^-}$ between the two chain maps by pairing with (1.17); the master equation (1.19) implies that:

$$\Phi(\hat{X}, \hat{\lambda}^1)_{\hat{J}^1,\theta^1} - \Phi(\hat{X}, \hat{\lambda}^0)_{\hat{J}^0,\theta^0} = d_{J^-,\theta^-} K(\hat{X}, \{d\hat{\lambda}^t\}_t)_{\{\hat{J}^t\}_t,\theta} + K(\hat{X}, \{d\hat{\lambda}^t\}_t)_{\{\hat{J}\}_t,\theta} d_{J^+,\theta^+} \quad (1.21)$$

Remark 1.2 (Higher homotopies associated to families of cobordisms). The proof of Theorem III generalizes immediately to the case of deformations parameterized by $t \in \Delta^n$. That is, there are sets $\Theta_{\mathrm{III}(n)} = \Theta_{\mathrm{III}(n)}(\hat{X}, (\hat{\omega}^t, \hat{J}^t)_{t\in\Delta^n})$ (specializing to $\Theta_{\mathrm{III}(0)} = \Theta_{\mathrm{II}}$ and $\Theta_{\mathrm{III}(1)} = \Theta_{\mathrm{III}}$) along with surjective maps:

$$\Theta_{\mathrm{III}(n)}(\hat{X}, (\hat{\omega}^t, \hat{J}^t)_{t\in\Delta^n}) \twoheadrightarrow \lim_{\Delta^k \subsetneq \Delta^n} \left[\Theta_{\mathrm{III}(k)}(\hat{X}, (\hat{\omega}^t, \hat{J}^t)_{t\in\Delta^k}) \to (\Theta^+_{\mathrm{I}} \times \Theta^-_{\mathrm{I}})\right] \quad (1.22)$$

(limit in the category of sets over $\Theta^+_{\mathrm{I}} \times \Theta^-_{\mathrm{I}}$). Furthermore, there are associated virtual moduli counts satisfying the natural master equation, thus giving rise to "higher homotopies" $K(\hat{X}, \{d\hat{\lambda}^t\}_{t\in\Delta^n})_{\{\hat{J}^t\}_{t\in\Delta^n},\theta}$ of degree n satisfying:

$$\begin{aligned} d_{J^-,\theta^-} K(\hat{X}, \{d\hat{\lambda}^t\}_{t\in\Delta^n})_{\{\hat{J}^t\}_{t\in\Delta^n},\theta} &+ (-1)^n K(\hat{X}, \{d\hat{\lambda}^t\}_{t\in\Delta^n})_{\{\hat{J}^t\}_{t\in\Delta^n},\theta} d_{J^+,\theta^+} \\ &= \sum_{\substack{k=0 \\ i:\Delta^{[0\ldots\hat{k}\ldots n]} \hookrightarrow \Delta^{[0\ldots n]}}}^{n} (-1)^k K(\hat{X}, \{d\hat{\lambda}^{i(t)}\}_{t\in\Delta^{n-1}})_{\{\hat{J}^{i(t)}\}_{t\in\Delta^{n-1}},i^*\theta} \end{aligned} \quad (1.23)$$

Remark 1.3 (Chain homotopy vs DGA homotopy). In Theorems III and IV, we obtain ordinary *chain homotopies* by counting *disconnected* curves. For certain purposes (e.g. to have good notions of augmentations and their linearized contact homology), one needs *dga homotopies* (in some suitable sense), and for this it seems to be necessary to count *connected* curves. Eliashberg–Givental–Hofer [EGH00, §2.4] sketch one way to count connected curves to obtain a dga homotopy (in one sense); see also Ekholm–Oancea [EO15, §§5.4–5.5]. In the context of Legendrian contact homology, Ekholm–Honda–Kálmán [EHK12, Lemma 3.13] give another way to count connected curves to obtain a dga homotopy (in another sense).

1.4 (IV) The composition homotopy

Setup IV. This setup consists of triples $(\hat{X}^{01}, \hat{\omega}^{01}, \hat{J}^{01})$, $(\hat{X}^{12}, \hat{\omega}^{12}, \hat{J}^{12})$ and triples (Y^0, λ^0, J^0), (Y^1, λ^1, J^1), (Y^2, λ^2, J^2). Here $(\hat{X}^{01}, \hat{\omega}^{01}, \hat{J}^{01})$, (Y^0, λ^0, J^0), (Y^1, λ^1, J^1) is as in Setup II, as is $(\hat{X}^{12}, \hat{\omega}^{12}, \hat{J}^{12})$, (Y^1, λ^1, J^1), (Y^2, λ^2, J^2).

Furthermore, we specify a smooth deformation of data $(\hat{X}^{02,t}, \hat{\omega}^{02,t}, \hat{J}^{02,t})_{t\in[0,\infty)}$ as in Setup II, with (Y^0, λ^0, J^0), (Y^2, λ^2, J^2) as positive/negative ends. For sufficiently large t, we identify $\hat{X}^{02,t}$ with the result of gluing $\hat{X}^{01} \sqcup \hat{X}^{12}$ by truncating the ends to $(-t, 0] \times Y_1 \subseteq \hat{X}^{01}$ and $[0, t) \times Y_1 \subseteq \hat{X}^{12}$ and identifying them via translation. Under this identification, $\hat{J}^{02,t}$ must coincide with the descent of $\hat{J}^{01} \sqcup \hat{J}^{12}$, and $\hat{\omega}^{02,t}$ must coincide with the descent of $\hat{\omega}^{01} \sqcup \hat{\omega}^{12}$ to $\hat{X}^{02,t}$ (only well-defined up to scale). If t is omitted, it means $t = 0$.

By considering appropriate moduli spaces $\overline{\mathcal{M}}_{\mathrm{IV}}(\{\gamma^+_i, \Gamma^-_i; \beta_i\}_{i\in I})$ of (possibly disconnected) $\hat{J}^{02,t}$-holomorphic curves in $\hat{X}^{02,t}$ for $t \in [0, \infty]$, we prove the following result.

Theorem IV. *Fix data as in Setup IV; let $\Theta^i_{\mathrm{I}} = \Theta_{\mathrm{I}}(Y^i, \lambda^i, J^i)$ and let $\Theta^{ij}_{\mathrm{II}} = \Theta_{\mathrm{II}}(\hat{X}^{ij}, \hat{\lambda}^{ij}, \hat{J}^{ij})$. There exists a set Θ_{IV} along with a surjective map $\Theta_{\mathrm{IV}} \twoheadrightarrow \Theta^{02}_{\mathrm{II}} \times_{\Theta^0_{\mathrm{I}} \times \Theta^2_{\mathrm{I}}} (\Theta^{01}_{\mathrm{II}} \times_{\Theta^1_{\mathrm{I}}} \Theta^{12}_{\mathrm{II}})$ and virtual moduli counts:*

$$\#\overline{\mathcal{M}}_{\mathrm{IV}}(\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I})^{\mathrm{vir}}_\theta \in \left(\bigotimes_{i\in I} \mathfrak{o}^\vee_{\gamma_i^+} \otimes \mathfrak{o}_{\Gamma_i^-}\right) \otimes \mathbb{Q} \tag{1.24}$$

This data is functorial in the data from Setup IV, and the following properties are satisfied for all $\theta \in \Theta_{\mathrm{IV}}$:

i. If $\sum_{i\in I} \mu(\gamma_i^+, \Gamma_i^-, \beta_i) \neq -1$, then $\#\overline{\mathcal{M}}_{\mathrm{IV}}(\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I})^{\mathrm{vir}}_\theta = 0$.

ii. If $\sum_{i\in I} \mu(\gamma_i^+, \Gamma_i^-, \beta_i) = -1$ and $\overline{\mathcal{M}}_{\mathrm{IV}}(\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I})$ is regular, then $\overline{\mathcal{M}}_{\mathrm{IV}}(\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I}) = \mathcal{M}_{\mathrm{IV}}(\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I})$ is a compact manifold of dimension zero naturally oriented by $\bigotimes_{i\in I} \mathfrak{o}_{\gamma_i^+} \otimes \mathfrak{o}^\vee_{\Gamma_i^-}$, and:

$$\#\overline{\mathcal{M}}_{\mathrm{IV}}(\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I})^{\mathrm{vir}}_\theta = \#\mathcal{M}_{\mathrm{IV}}(\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I}) \tag{1.25}$$

iii. The numbers $\#\overline{\mathcal{M}}_{\mathrm{IV}}(\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I})^{\mathrm{vir}}_\theta$ satisfy the following "master equation":

$$\#\partial\overline{\mathcal{M}}_{\mathrm{IV}}(\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I})^{\mathrm{vir}}_\theta = 0 \tag{1.26}$$

for all $\{(\gamma_i^+, \Gamma_i^-, \beta_i)\}_{i\in I}$, where the left hand side denotes the sum over all codimension one boundary strata of the relevant products of (1.2) (using $\theta^{0,1,2}$), (1.11) (using $\theta^{01,12,02}$), (1.24), and inverse covering multiplicities of intermediate orbits (this sum is finite by compactness).

For data as in Setup IV, it follows that the following diagram commutes:

$$\begin{tikzcd}
 & CH_\bullet(Y^1, \xi^1)_{\lambda^1, J^1, \theta^1} \arrow[dr, "\Phi(\hat{X}^{12}, d\hat{\lambda}^{12})_{\hat{J}^{12}, \theta^{12}}"] & \\
CH_\bullet(Y^0, \xi^0)_{\lambda^0, J^0, \theta^0} \arrow[ur, "\Phi(\hat{X}^{01}, d\hat{\lambda}^{01})_{\hat{J}^{01}, \theta^{01}}"] \arrow[rr, "\Phi(\hat{X}^{02}, d\hat{\lambda}^{02})_{\hat{J}^{02}, \theta^{02}}"'] & & CH_\bullet(Y^2, \xi^2)_{\lambda^2, J^2, \theta^2}
\end{tikzcd} \tag{1.27}$$

for $(\theta^{0,1,2}, \theta^{01,12,02}) \in \Theta^{02}_{\mathrm{II}} \times_{\Theta^0_{\mathrm{I}} \times \Theta^2_{\mathrm{I}}} (\Theta^{01}_{\mathrm{II}} \times_{\Theta^1_{\mathrm{I}}} \Theta^{12}_{\mathrm{II}})$ (one obtains a chain homotopy between the two chain maps by pairing with (1.24) and using (1.26)).

1.5 The contact homology functor

We now assemble our main results to define the contact homology functor (1.28).

Let $(\mathfrak{Contact}, \mathfrak{Exact})_n$ denote the category whose objects are closed co-oriented contact manifolds (Y^{2n-1}, ξ) and whose morphisms are deformation classes of exact symplectic cobordisms (X^{2n}, λ). Let $\mathfrak{Ring}^{\mathbb{Z}/2}_{\mathbb{Q}}$ denote the category whose objects are supercommutative $\mathbb{Z}/2$-graded unital $\mathbb{Q}$-algebras and whose morphisms are graded unital $\mathbb{Q}$-algebra homomorphisms. Contact homology is a symmetric monoidal functor:

$$CH_\bullet : (\mathfrak{Contact}, \mathfrak{Exact})^{\sqcup}_n \to (\mathfrak{Ring}^{\mathbb{Z}/2}_{\mathbb{Q}})^{\otimes} \tag{1.28}$$

The symmetric monoidal structure on $(\mathfrak{Contact}, \mathfrak{Exact})_n$ is disjoint union $\sqcup$, and the symmetric monoidal structure on $\mathfrak{Ring}_{\mathbb{Q}}^{\mathbb{Z}/2}$ is the super tensor product $\otimes$ ($A \otimes B$ is endowed with the multiplication $(a \otimes b)(a' \otimes b') := (-1)^{|a'||b|} aa' \otimes bb'$, and the isomorphism $A \otimes B \xrightarrow{\sim} B \otimes A$ is given by $a \otimes b \mapsto (-1)^{|a||b|} b \otimes a$).

We now construct the functor (1.28) from Theorems I, II, III, IV in detail. Concretely, this means we should define:

- For every co-oriented contact manifold (Y, ξ), a supercommutative $\mathbb{Z}/2$-graded unital $\mathbb{Q}$-algebra $CH_\bullet(Y, \xi)$.
- For every exact symplectic cobordism (X, λ) from (Y^+, ξ^+) to (Y^-, ξ^-), a graded unital $\mathbb{Q}$-algebra map $\Phi(X, \lambda) : CH_\bullet(Y^+, \xi^+) \to CH_\bullet(Y^-, \xi^-)$.
- Isomorphisms $CH_\bullet(M, \xi) \otimes CH_\bullet(M', \xi') = CH_\bullet(M \sqcup M', \xi \sqcup \xi')$.

such that:

- The morphism associated to the identity cobordism is the identity map.
- The morphism $\Phi(X^{02}, \lambda^{02})$ associated to a composition of exact symplectic cobordisms $X^{02} = X^{01} \# X^{12}$ coincides with the composition $\Phi(X^{12}, \lambda^{12}) \circ \Phi(X^{01}, \lambda^{01})$.
- The morphism $\Phi(X, \lambda)$ depends only on the deformation class of (X, λ).
- The isomorphisms $CH_\bullet(M, \xi) \otimes CH_\bullet(M', \xi') = CH_\bullet(M \sqcup M', \xi \sqcup \xi')$ are commutative, associative, and compatible with the cobordism maps.

The construction is as follows.

Theorem I provides a supercommutative $\mathbb{Z}/2$-graded unital $\mathbb{Q}$-algebra:

$$CH_\bullet(Y, \xi)_{\lambda, J, \theta} \tag{1.29}$$

for any co-oriented contact manifold (Y, ξ) with non-degenerate contact form λ, admissible almost complex structure J, and $\theta \in \Theta_{\mathrm{I}}(Y, \lambda, J)$.

Theorem II provides a graded unital $\mathbb{Q}$-algebra map:

$$CH_\bullet(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \xrightarrow{\Phi(\hat{X}, \hat{\lambda})_{\hat{J}, \theta}} CH_\bullet(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-} \tag{1.30}$$

for any exact symplectic cobordism (X, λ) with $\lambda^\pm$ non-degenerate, admissible almost complex structure $\hat{J}$ coinciding with $\hat{J}^\pm$ near infinity, and $\theta \in \Theta_{\mathrm{II}}(\hat{X}, \hat{\lambda}, \hat{J})$ mapping to $\theta^\pm \in \Theta_{\mathrm{I}}^\pm$.

Theorem III shows that the following two maps coincide:

$$CH_\bullet(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \underset{\Phi\hat{X}, \hat{\lambda}^1)_{\hat{J}^1, \theta^1}}{\overset{\Phi(\hat{X}, \hat{\lambda}^0)_{\hat{J}^0, \theta^0}}{\rightrightarrows}} CH_\bullet(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-} \tag{1.31}$$

Note that this immediately implies that $\Phi(\hat{X}, \hat{\lambda})_{\hat{J}, \theta}$ is independent of $\hat{J}$ and θ, and depends only on the deformation class of (X, λ). Thus we may rewrite (1.30) as:

$$CH_\bullet(Y^+, \xi^+)_{\lambda^+, J^+, \theta^+} \xrightarrow{\Phi(X, \lambda)} CH_\bullet(Y^-, \xi^-)_{\lambda^-, J^-, \theta^-} \tag{1.32}$$

Theorem IV shows that the following diagram commutes:

$$
\begin{tikzcd}
 & CH_\bullet(Y^1,\xi^1)_{\lambda^1,J^1,\theta^1} \arrow[dr, "\Phi(X^{12},\lambda^{12})"] & \\
CH_\bullet(Y^0,\xi^0)_{\lambda^0,J^0,\theta^0} \arrow[ur, "\Phi(X^{01},\lambda^{01})"] \arrow[rr, "\Phi(X^{02},\lambda^{02})"'] & & CH_\bullet(Y^2,\xi^2)_{\lambda^2,J^2,\theta^2}
\end{tikzcd}
\tag{1.33}
$$

Lemma 1.4. *Let (Y,ξ) be a co-oriented contact manifold with two non-degenerate contact forms λ^+,λ^-. Let $(\hat{X},\hat{\lambda})$ denote the trivial exact symplectic cobordism from (Y,λ^+) to (Y,λ^-). Then the map:*

$$CH_\bullet(Y,\xi)_{\lambda^+,J^+,\theta^+} \xrightarrow{\Phi(\hat{X},\hat{\lambda})} CH_\bullet(Y,\xi)_{\lambda^-,J^-,\theta^-} \tag{1.34}$$

is an isomorphism for any $J^\pm$ and $\theta^\pm$.

Proof. In view of the commutativity of (1.33), it suffices to treat the case $\lambda^+=\lambda^-=\lambda$ and $J^+=J^-=J$.

Choose the $\mathbb{R}$-invariant almost complex structure $\hat{J}=\hat{J}^\pm$ on $\hat{X}$, and choose any $\theta\in\Theta_{\mathrm{II}}$ mapping to $\theta^\pm$. We will show that the map on chains:

$$CC_\bullet(Y,\xi)_{\lambda,J,\theta^+} \xrightarrow{\Phi(\hat{X},\hat{\lambda})_{\hat{J},\theta}} CC_\bullet(Y,\xi)_{\lambda,J,\theta^-} \tag{1.35}$$

is an isomorphism, which is clearly sufficient.

We consider the ascending filtration on both sides of (1.35) whose $^{\leq(a,k)}$ filtered piece is the $\mathbb{Q}$-subspace generated by all monomials of Reeb orbits with total action $< a$ or total action $= a$ and degree $\geq k$. The map (1.35) respects this filtration. Indeed, the integral of $d\lambda$ (here we consider $d\lambda$ and not $d\hat{\lambda}$) over any $\hat{J}$-holomorphic curve is ≥ 0, with equality iff the curve is a branched cover of a trivial cylinder, and every branched cover of a trivial cylinder has at least one negative end. Since the filtration is well-ordered, it suffices to show that the induced map on associated gradeds is an isomorphism.

The curves contributing to the action of (1.35) on associated gradeds are the branched covers of trivial cylinders with exactly one negative end, and such curves are themselves necessarily trivial cylinders by Riemann–Hurwitz. Since there is exactly one such trivial cylinder for every Reeb orbit, it suffices to show that trivial cylinders are cut out transversally. This is a standard fact, whose proof we recall in Lemma 2.24. $\square$

Now for a contact manifold (Y,ξ), all groups $CH_\bullet(Y,\xi)_{\lambda,J,\theta}$ are canonically isomorphic via the morphisms $\Phi(\hat{X},\hat{\lambda})$ associated to the trivial cobordisms (by Lemma 1.4 and the commutativity of (1.33)). Thus we get a well-defined object:

$$CH_\bullet(Y,\xi) \tag{1.36}$$

independent of λ, J, θ. Formally speaking, $CH_\bullet(Y,\xi)$ is the limit (and the colimit) of $\{CH_\bullet(Y,\xi)_{\lambda,J,\theta}\}_{\lambda,J,\theta}$, which is attained at any particular triple (λ,J,θ). Note that for any contact structure, the set of non-degenerate contact forms is generic (and in particular non-empty).

The commutativity of (1.33) also implies that a deformation class of exact symplectic coboordism (X,λ) from (Y^+,ξ^+) to (Y^-,ξ^-) induces a well-defined graded unital $\mathbb{Q}$-algebra map:

$$\Phi(X,\lambda) : CH_\bullet(Y^+,\xi^+) \to CH_\bullet(Y^-,\xi^-) \tag{1.37}$$

and that $\Phi(X^{02},\lambda^{02}) = \Phi(X^{12},\lambda^{12}) \circ \Phi(X^{01},\lambda^{01})$ for $X^{02} = X^{01}\#X^{12}$.

To construct the symmetric monoidal structure on $CH_\bullet$, it suffices to observe that the sets Θ_{I}, Θ_{II} are themselves (almost) symmetric monoidal, in the sense made precise in Proposition 4.39. This completes the construction of the contact homology functor (1.28) in terms of the main results Theorems I, II, III, IV.

1.6 Variations on contact homology

We now recall (following [Eli98] and [EGH00]) some important variations on the basic contact invariant $CH_\bullet(Y,\xi)$ defined above.

- (Grading by $H_1(Y)$) Contact homology $CH_\bullet(Y,\xi)$ has a grading by $H_1(Y)$ (the grading of a given monomial in Reeb orbits equals its total homology class).

- (Refinement of $\mathbb{Z}/2$-grading) Contact homology $CH_\bullet(Y,\xi)$ has a relative grading by $\mathbb{Z}/2c_1(\xi)\cdot H_2(Y)$, which is absolute over the $0 \in H_1(Y)$ graded piece. The grading is given by $|\gamma| = \mathrm{CZ}(\gamma) + n - 3$.

- (Action filtration) If we equip (Y,ξ) with a contact form λ, then for $a \in \mathbb{R}$, there is an invariant $CH_\bullet(Y,\lambda)^{<a}$ equipped with functorial maps $CH_\bullet(Y,\lambda)^{<a} \to CH_\bullet(Y,\lambda')^{<a'}$ for $\frac{\lambda}{a} \geq \frac{\lambda'}{a'}$ such that:
$$CH_\bullet(Y,\xi) = \varinjlim CH_\bullet(Y,\lambda)^{<a} \tag{1.38}$$
Namely, $CH_\bullet(Y,\lambda)^{<a}$ is defined as the homology $CC_\bullet(Y,\xi)^{<a}_{\lambda,J,\theta} \subseteq CC_\bullet(Y,\xi)_{\lambda,J,\theta}$, the subspace spanned by monomials of total action $< a$ (note that the differential strictly decreases action). This invariant $CH_\bullet^{<a}$ may be constructed out of Theorems I, II, III, IV as in §1.5.

- (Coefficients in $\mathbb{Q}[H_2(Y)]$) Contact homology $CH_\bullet(Y,\xi)$ has a natural lift $\overline{CH}_\bullet(Y,\xi)$ to the group ring $\mathbb{Q}[H_2(Y;\mathbb{Z})]$. More intrinsically, $\overline{CH}_\bullet$ may be thought of as a local system over the space of 1-cycles in Y, namely $\tau_{\geq 0}C_{\bullet+1}(Y)$. Contact homology with group ring coefficients $\overline{CH}_\bullet(Y,\xi)$ has a relative $\mathbb{Z}$-grading, where $\mathbb{Q}[H_2(Y)]$ is $\mathbb{Z}$-graded by $2c_1(\xi) : H_2(Y) \to \mathbb{Z}$. This invariant $\overline{CH}_\bullet$ may be constructed out of Theorems I, II, III, IV as in §1.5.

- (Contact homology of contractible orbits) There is an invariant $CH_\bullet^{\mathrm{contr}}(Y,\xi)$ (an algebra) obtained from the chain complex $CC_\bullet^{\mathrm{contr}}(Y,\lambda)$ generated as an algebra by contractible Reeb orbits (with a differential which counts curves whose asymptotic orbits are all contractible). There is also an invariant $CH_\bullet^{\alpha}(Y,\xi)$ (a module over $CH_\bullet^{\mathrm{contr}}(Y,\xi)$) obtained from the chain complex $CC_\bullet^{\alpha}(Y,\lambda)$ generated as a module over $CC_\bullet^{\mathrm{contr}}(Y,\lambda)$ by Reeb orbits in a fixed nontrivial homotopy class α (with differential counting curves whose asymptotic orbits are either all contractible or all contractible except for the

positive end and one negative end both in class α). These invariants $CH_\bullet^{\mathrm{contr}}$ and $CH_\bullet^\alpha$ may be constructed out of Theorems I, II, III, IV as in §1.5.

- (Cylindrical contact homology) If (Y,ξ) is *hypertight* (admits a contact form with no contractible Reeb orbits) then there is an invariant $CH_\bullet^{\mathrm{cyl}}(Y,\xi)$ defined as follows. If (Y,ξ) admits a *non-degenerate* contact form with no contractible Reeb orbits, then $CH_\bullet^{\mathrm{cyl}}(Y,\xi)$ is defined as the homology of the complex $CC_\bullet^{\mathrm{cyl}}(Y,\lambda) := \bigoplus_{\gamma\in\mathcal{P}_{\mathrm{good}}} \mathfrak{o}_\gamma$ with the differential which counts pseudo-holomorphic cylinders. If this is not the case, then one must first define $CH_\bullet^{\mathrm{cyl}}(Y,\lambda)^{<a}$ for non-degenerate contact forms λ with no contractible Reeb orbits of action $< a$, and then let $CH_\bullet^{\mathrm{cyl}}(Y,\xi) := \varinjlim CH_\bullet^{\mathrm{cyl}}(Y,\lambda)^{<a}$. This invariant $CH_\bullet^{\mathrm{cyl}}$ may be constructed out of Theorems I, II, III, IV as in §1.5.

- (Linearized contact homology) Given a contact manifold (Y,ξ) equipped with an exact symplectic filling (X,λ), it should be possible to define an invariant $CH_\bullet^{\mathrm{lin}}(X,\lambda)$ as the homology of $CC_\bullet^{\mathrm{lin}}(X,\lambda) := \bigoplus_{\gamma\in\mathcal{P}_{\mathrm{good}}} \mathfrak{o}_\gamma$ with differential which counts pseudo-holomorphic buildings in $\hat{Y}$ "anchored" in $\hat{X}$ as in [BEE12]. It should be straightforward to generalize the methods of this paper to construct this invariant.

- (Invariants of contactomorphisms) There is a natural homomorphism:

$$\pi_0 \operatorname{Cont}(Y,\xi) \to \operatorname{Aut}_{\mathbb{Q}}(CH_\bullet(Y,\xi)) \tag{1.39}$$

  namely the tautological action of $\operatorname{Cont}(Y,\xi)$ on $CH_\bullet(Y,\xi)$. This action admits the following description in terms of cobordism maps which shows that it descends to π_0. For any $\varphi \in \operatorname{Cont}(Y,\xi)$, denote by X_φ the exact symplectic cobordism from Y to itself obtained from the trivial cobordism by changing the marking on the negative end by φ. The action of φ on $CH_\bullet(Y,\xi)$ clearly coincides with the cobordism map $\Phi(X_\varphi)$. On the other hand, $\Phi(X_\varphi)$ only depends on the class of φ in π_0 since cobordism maps are invariant under deformation. It is also clear that $\varphi \mapsto \Phi(X_\varphi)$ is a group homomorphism since $X_\varphi \# X_\psi = X_{\varphi\psi}$.

  As pointed out by P. Massot, any contactomorphism *symplectically pseudo-isotopic* to the identity (a notion due to Cieliebak–Eliashberg [CE12]) lies in the kernel of (1.39); indeed, φ is symplectically pseudo-isotopic to the identity iff X_φ is isomorphic to X_{id} as symplectic cobordisms from Y to itself. Thus a contactomorphism which acts non-trivially on contact homology cannot be symplectically pseudo-isotopic to the identity.

  The above construction generalizes to give a natural homomorphism:

$$H_k(\operatorname{Cont}(Y,\xi)) \to \operatorname{Hom}_{\mathbb{Q}}(CH_\bullet(Y,\xi), CH_{\bullet+k}(Y,\xi)) \tag{1.40}$$

  Namely, a family of $\varphi \in \operatorname{Cont}(Y,\xi)$ gives a family of cobordisms X_φ, and the "higher homotopies" defined in Remark 1.2 give the desired map (1.40). For non-trivial examples, we refer the reader to Bourgeois [Bou06], who introduced (1.39)–(1.40) (or, rather, their pre-compositions with the map $\Omega_\xi\Xi(Y) \to \operatorname{Cont}(Y,\xi)$ coming from Gray's fibration sequence $\operatorname{Cont}(Y,\xi) \to \operatorname{Diff}(Y,\xi) \to \Xi(Y)$, where $\Xi(Y)$ denotes the space of contact structures on Y and Ω_ξ denotes the space of loops based at $\xi \in \Xi(Y)$).

It should be possible to upgrade (1.39)–(1.40) into the statement that $CC_\bullet(Y,\xi)$ is an A_∞-module over $C_\bullet(\operatorname{Cont}(Y,\xi))$ (i.e. contact chains of (Y,ξ) should be a "derived local system" over $B\operatorname{Cont}(Y,\xi)$). One way to approach the construction of such a structure would be to simultaneously generalize Remark 1.2 and Theorem IV to compositions of multiple cobordisms. It should be straightforward to generalize the methods of this paper to carry out this argument.

- (Integer coefficients) An interesting open question (promoted by Abouzaid) is how to naturally lift contact homology from $\mathbb{Q}$ to $\mathbb{Z}$.

1.7 Applications and calculations

We now mention a few applications and calculations of contact homology.

1.7.1 Overtwisted contact manifolds

A given (connected, non-empty) contact manifold is either *tight* or *overtwisted*. Overtwisted contact structures are classified completely by an h-principle due to Eliashberg [Eli89] in dimension three and Borman–Eliashberg–Murphy [BEM14] in general.

Contact homology (even with group ring coefficients) vanishes on any overtwisted contact manifold. In dimension three, this is a result of Eliashberg [Eli98, p334, Theorem 3.5(2)] (a proof is given in Yau [Yau06] and the appendix by Eliashberg). In higher dimensions, this follows from the result of Bourgeois–vanKoert [BvK10, Theorem 1.3] that contact homology vanishes for any contact manifold admitting a negatively stabilized open book, together with the result of Casals–Murphy–Presas [CMP15, Theorem 1.1] that a contact manifold admits a negatively stabilized open book iff it is overtwisted. These vanishing results are proved by exhibiting a contact form with a non-degenerate Reeb orbit bounding exactly one pseudo-holomorphic plane in the symplectization (which is cut out transversally); in particular, they are valid for the contact homology we construct here.

Contact homology (with group ring coefficients) should also vanish for PS-overtwisted contact manifolds. The notion of PS-overtwistedness is due to Niederkrüger [Nie06] and Massot–Niederkrüger–Wendl [MNW13]; every overtwisted contact manifold is PS-overtwisted (see [BEM14, p4]), and the converse is currently open (except in dimension three where it holds by definition). The argument for vanishing (due to Bourgeois–Niederkrüger and sketched in [Bou09]) considers (virtual) counts for (compactified) moduli spaces of pseudo-holomorphic disks with boundary on the plastikstufe (or bLob), one marked point mapping to a fixed curve on the plastikstufe from its core to its boundary, and an arbitrary number of negative punctures. It should be straightforward to generalize the methods of this paper to carry out this argument.

It is natural to ask whether contact homology detects overtwistedness, that is, whether $CH_\bullet(Y,\xi)=0$ implies that (Y,ξ) is overtwisted (this would imply that Legendrian surgery preserves tightness, recently proven in dimension three by Wand [Wan14]). Recent examples of Ekholm [Ekh15] suggest there may be counterexamples, at least in higher dimensions. The somewhat similar Heegaard Floer contact invariant $c(\xi)\in\widehat{HF}(-Y)/\{\pm 1\}$ of Ozsváth–Szabó [OS05] also vanishes for overtwisted contact manifolds, and is known to not detect

overtwistedness by examples of Ghiggini [Ghi06].

1.7.2 Existence of Reeb orbits

The *Weinstein conjecture* [Wei79] asserts that every contact form on a closed manifold admits at least one Reeb orbit. In dimension three, a breakthrough was made by Hofer [Hof93], who proved the Weinstein conjecture for contact three-manifolds which are either S^3, overtwisted, or have nontrivial π_2 (using pseudo-holomorphic curves in symplectizations). Later, Taubes [Tau07] proved the Weinstein conjecture for all contact three-manifolds (using Seiberg–Witten Floer homology and Embedded Contact Homology of Hutchings and Hutchings–Taubes [Hut02, Hut09, HT07, HT09a]). In higher dimensions, a number of cases are known, for example [AH09, NR11, GZ12, DGZ14, GZ14, CDvK14].

If $CH_\bullet(Y,\xi)\neq\mathbb{Q}$, then the Weinstein conjecture holds for (Y,ξ). Indeed, the existence of a Reeb orbit for non-degenerate λ is immediate. Moreover, given any fixed (λ_0,a_0) such that $CH_\bullet(Y,\lambda_0)^{<a_0}\to CH_\bullet(Y,\xi)$ does not factor through $\mathbb{Q}\to CH_\bullet(Y,\xi)$, every non-degenerate λ has a Reeb orbit of action $\leq a_0\sup\frac{\lambda}{\lambda_0}$. Now the same holds for arbitrary contact forms λ since non-degenerate contact forms are generic (and hence dense).

It is natural to ask whether $CH_\bullet(Y,\xi)=\mathbb{Q}$ for any (non-empty) contact manifold.

One can also show the existence of Reeb orbits in particular homology or homotopy classes by taking advantage of the grading of $CH_\bullet(Y,\xi)$ by $H_1(Y)$ or by using $CH^\alpha_\bullet(Y,\xi)$. In particular, if $CH_\bullet(Y,\xi)=0$, then every contact form for ξ has a *contractible* Reeb orbit. Note that it is already known by work of Albers–Hofer [AH09] that PS-overtwisted contact manifolds always have a contractible Reeb orbit.

One can also use contact homology to estimate the growth rate of the number of Reeb orbits below a given action threshold (see, e.g. Vaugon [Vau15]). Many such results have been obtained using Embedded Contact Homology, e.g. [HT09b, CGHR15, CGH14], and there are also earlier results due to Hofer–Wysocki–Zehnder [HWZ03]. Ginzburg–Kerman [GK10] used contact homology to give other interesting restrictions on Reeb dynamics.

1.7.3 Fillability and capacities

A contact manifold with vanishing contact homology is not exactly symplectically fillable; more generally, if the positive end of an exact symplectic cobordism has vanishing contact homology then so does the negative end (since there are no unital ring maps $0\to R$ for a non-zero ring R).

In fact, a contact manifold with vanishing contact homology is not strongly symplectically fillable, as shown by Niederkrüger–Wendl [NW11, Corollary 6(1)] (stated in dimension three, though the proof works in general). Moreover, the vanishing of contact homology with certain group ring coefficients can obstruct weak fillability as well, as shown by Niederkrüger–Wendl [NW11, Theorem 6, Corollary 6(2)] in dimension three and Massot–Niederkrüger–Wendl [MNW13, Corollary 7, Theorem F] in general. Our results provide sufficient virtual curve counts for both of these results. These filling obstructions have been generalized by Latschev–Wendl [LW11] assuming the existence of certain SFT invariants.

Symplectic embedding capacities can be defined from the action filtration on contact homology; they are expected to coincide with the capacities defined by Ekeland–Hofer [EH89,

EH90]. Similar (though different) capacities have been defined by Hutchings [Hut11] using Embedded Contact Homology, see also [Hut15].

1.7.4 Contact non-squeezing

Eliashberg–Kim–Polterovich [EKP06] proved certain contact non-squeezing results (in similar spirit to Gromov's non-squeezing theorem) using contact homology. In their setting, generic transversality holds, and our results are not needed.

1.7.5 Relation to positive S^1-equivariant symplectic homology

The linearized contact homology of (X, λ) should coincide with $SH_\bullet^{+,S^1}(X, \lambda) \otimes_{\mathbb{Z}} \mathbb{Q}$, where $SH_\bullet^{+,S^1}$ denotes S^1-*equivariant positive symplectic homology* (see [BO14]). This isomorphism was proven by Bourgeois–Oancea [BO09, BO15] under certain transversality assumptions. It should be possible to generalize the methods of this paper to prove this isomorphism in general. One can use $SH_\bullet^{+,S^1}$ in place of contact homology for some applications (see Gutt [Gut15]).

1.7.6 Calculations

It is trivial to calculate contact homology given a non-degenerate contact form all of whose Reeb orbits are even, since the differential then vanishes for index reasons. Even this simple case yields interesting results. Ustilovsky [Ust99] constructed such "even contact forms" on certain Brieskorn spheres, leading to the conclusion that S^{4k+1} carries infinitely many non-contactomorphic tight contact structures, in every homotopy class of almost contact structure. Abreu–Macarini [AM12] constructed even contact forms on any "good toric contact manifold", leading to the conclusion that $S^2 \times S^3$ carries infinitely many tight contact structures in the (unique) homotopy class of almost contact structures with $c_1 = 0$.

It is also sometimes possible to analyze pseudo-holomorphic curves directly to calculate contact homology. Examples of such calculations (together with applications) include [EGH00, Theorem 1.9.9], [BC05], [CH13], [vK08], [Yau04] (though see relevant remarks in [BO15]), [Yau09]. Calculations are made much more convenient if one is allowed to use Morse–Bott contact forms as in Bourgeois [Bou02] (see Remark 1.7).

The linearized contact homology of cotangent bundles should be given by $CH_\bullet^{\mathrm{lin}}(T^*N) = H_\bullet^{S^1}(LN, N)$ (S^1-equivariant homology of the loop space LN relative to the constant loops, possibly with twisted coefficients as in Abouzaid [Abo14]). This is originally due to Cieliebak–Latschev [CL09], and is consistent with the expectation $CH_\bullet^{\mathrm{lin}}(T^*N) = SH_\bullet^{+,S^1}(T^*N) \otimes_{\mathbb{Z}} \mathbb{Q}$ and Viterbo's theorem [Abo14].

It is natural to investigate the effect of Weinstein surgery on contact homology. In particular there should be a long exact sequence involving $CH_\bullet(Y_\Lambda) \to CH_\bullet(Y)$ for any isotropic sphere $\Lambda \subseteq Y$. For Legendrian surgery, Bourgeois–Ekholm–Eliashberg [BEE12] sketch the construction of such a long exact sequence in linearized contact homology. For subcritical surgery, work of Yau [Yau04] (see relevant remarks in [BO15]) is closely related.

1.8 Remarks for the experts

We collect here some remarks for the experts concerning various technical aspects of our approach.

Remark 1.5 (Small compactifications). We choose to use compactifications of the relevant moduli spaces which are smaller than the "standard SFT compactifications" used in [EGH00, BEHWZ03]. Roughly speaking, we collapse all trivial cylinders, and we do not keep track of the relative vertical position of different components of disconnected curves in symplectizations $\mathbb{R} \times Y$.

Our alternative compactifications are more convenient for proving the master equations of contact homology: the codimension one boundary strata in our compactifications correspond bijectively with the desired terms in the "master equations", whereas the usual SFT compactifications contain additional codimension one boundary strata. If we were to use the usual SFT compactifications, we would need to additionally argue that the contribution of each such extra codimension one boundary stratum vanishes.

Remark 1.6 (Topological gluing theorem). We prove a gluing theorem which gives a local topological description of our compactified moduli spaces over the locus where they are transverse. In order to apply the virtual fundamental cycle machinery developed in [Par15], a smooth structure on the compactified moduli spaces is not needed, nor is any gluing in non-transverse settings.

Remark 1.7 (Morse–Bott contact homology). It should be possible to define contact homology in terms of Morse–Bott contact forms, by counting appropriate pseudo-holomorphic cascades (this is due to Bourgeois [Bou02]). It should be possible to generalize the methods of this paper to construct the relevant virtual moduli counts (though the relevant gluing analysis would be more technical).

Remark 1.8 (Symplectic Field Theory). It should be straightforward to generalize the methods of this paper to construct many of the more general Symplectic Field Theory invariants from [EGH00].

2 Moduli spaces of pseudo-holomorphic curves

In this section, we define the moduli spaces of pseudo-holomorphic curves which we will use to define contact homology.

2.1 Categories of strata $\mathcal{S}_{\mathrm{I}}$, $\mathcal{S}_{\mathrm{II}}$, $\mathcal{S}_{\mathrm{III}}$, $\mathcal{S}_{\mathrm{IV}}$

We begin by introducing collections $\mathcal{S}_{\mathrm{I}}$, $\mathcal{S}_{\mathrm{II}}$, $\mathcal{S}_{\mathrm{III}}$, $\mathcal{S}_{\mathrm{IV}}$ of labelled trees which we will use to index the strata of the compactified moduli spaces of pseudo-holomorphic curves. A labelled tree describes the "combinatorial type" of a pseudo-holomorphic curve: the tree is the dual graph of the domain, and it is labeled with the homotopy class and asymptotics of the map.

Gluing pseudo-holomorphic curves corresponds to contracting a subset of the edges of a tree and updating the labels accordingly. We regard $\mathcal{S}_{\mathrm{I}}$, $\mathcal{S}_{\mathrm{II}}$, $\mathcal{S}_{\mathrm{III}}$, $\mathcal{S}_{\mathrm{IV}}$ as categories, with such edge contractions as morphisms. Morphisms of labelled trees then correspond to inclusions of (closed) strata in compactified moduli spaces.

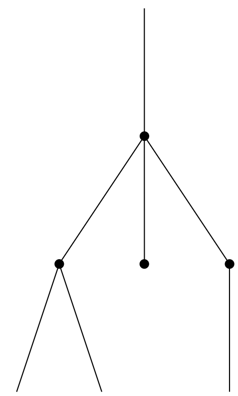

Figure 1: A tree. The edges are all directed downwards.

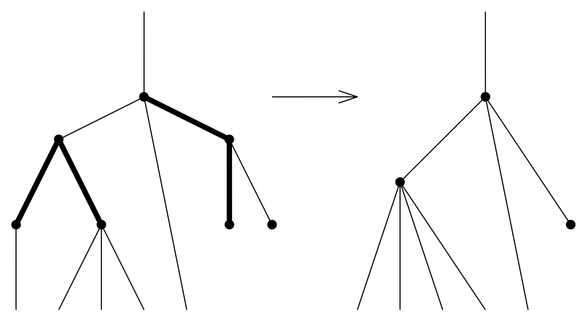

Figure 2: A contraction of trees. The edges which have been contracted are marked bold.

The categories $\mathcal{S}_{\mathrm{I}}$, $\mathcal{S}_{\mathrm{II}}$, $\mathcal{S}_{\mathrm{III}}$, $\mathcal{S}_{\mathrm{IV}}$ carry some additional (vaguely monoidal) structure, corresponding to the fact that boundary strata in compactified moduli spaces can be expressed as products of other "smaller" moduli spaces. The relevant operation on trees is that of "concatenation", where we take some collection of trees and identify some pairs of input/output edges with matching labels to produce a new tree.

Definition 2.1 (Tree)**.** A *tree* shall mean a finite directed tree, allowing "half edges" (i.e. edges with missing source or missing sink), in which every vertex has a unique incoming edge (see Figure 1). We denote by $(V(T), E(T))$ the sets of vertices and edges of a tree T. An incoming half edge is called an *input edge*, an outgoing half edge is called an *output edge*, and full edges are called *interior edges*. We write $E(T) = E^{\mathrm{int}}(T) \sqcup E^{+}(T) \sqcup E^{-}(T)$ for the partition into interior, input, and output edges (respectively), and we let $E^{\pm}(T) := E^{+}(T) \sqcup E^{-}(T)$. For a vertex $v \in V(T)$, we denote by $e^{+}(v) \in E(T)$ the unique incoming edge at v, and we denote by $\{e^{-}(v) \in E(T)\}$ the outgoing edges at v.

A *contraction* of trees $T \to T'$ shall mean a map obtained by contracting some collection of interior edges of T (see Figure 2). Every tree has a unique maximal contraction, namely contracting all interior edges.

Definition 2.2 (Category of strata $\mathcal{S}_{\mathrm{I}}$)**.** We define a category $\mathcal{S}_{\mathrm{I}}$ (depending on data as in Setup I). An object of $\mathcal{S}_{\mathrm{I}}$ is a connected non-empty tree T along with decorations consisting of a Reeb orbit $\gamma_e \in \mathcal{P}$ for all edges $e \in E(T)$, a homotopy class $\beta_v \in \pi_2(Y, \gamma_{e^{+}(v)} \sqcup \{\gamma_{e^{-}(v)}\})$

for all vertices $v \in V(T)$, and a basepoint $p_e \in \gamma_e$ (meaning a point on the underlying simple orbit of γ_e) for all input/output edges $e \in E^\pm(T)$.

A morphism $T \to T'$ in $\mathcal{S}_{\mathrm{I}}$ consists of a contraction of underlying trees, compatible with the decorations, along with paths between the basepoints. "Compatible with the decorations" means that $\gamma_{\pi(e)} = \gamma_e$ for all non-contracted edges $e \in E(T)$ and $\beta_{v'} = \#_{\pi(v)=v'}\beta_v$ for all vertices $v' \in V(T')$. A "path between basepoints" $p, p' \in \gamma \in \mathcal{P}$ means a homotopy class of paths in the underlying simple orbit of γ, modulo the relation that identifies two paths iff their "difference" lifts to γ (i.e. has degree divisible by the covering multiplicity d_γ).[6] Note that $\mathrm{Aut}(T)$, the automorphism group of an object $T \in \mathcal{S}_{\mathrm{I}}$, is the semi-direct product of $\prod_{e\in E^\pm(T)} \mathbb{Z}/d_{\gamma_e}$ with the subgroup of the automorphism group of the underlying tree preserving the decorations.

A *concatenation* $\{T_i\}_i$ *in* $\mathcal{S}_{\mathrm{I}}$ shall mean a finite collection of trees $T_i \in \mathcal{S}_{\mathrm{I}}$ along with a matching between some pairs of input/output edges of the T_i's (with matching γ_e) such that the resulting gluing is a connected tree, along with a choice of paths between the basepoints for each pair of matched edges. Given a concatenation $\{T_i\}_i$ in $\mathcal{S}_{\mathrm{I}}$, there is a resulting object $\#_i T_i \in \mathcal{S}_{\mathrm{I}}$. If $\{T_i\}_i$ is a concatenation in $\mathcal{S}_{\mathrm{I}}$ and $T_i = \#_j T_{ij}$ for some concatenations $\{T_{ij}\}_j$, there is a resulting *composite concatenation* $\{T_{ij}\}_{ij}$ with natural isomorphisms $\#_{ij}T_{ij} = \#_i\#_j T_{ij} = \#_i T_i$.

Call an object $T \in \mathcal{S}_{\mathrm{I}}$ *maximal* iff the only morphism out of T is the identity map (i.e. iff T has exactly one vertex). Every object $T \in \mathcal{S}_{\mathrm{I}}$ has a unique morphism to a maximal object, which is unique up to unqiue isomorphism. Note the following key fact: an object $T \in \mathcal{S}_{\mathrm{I}}$ is maximal iff it cannot be expressed nontrivially as a concatenation.

Definition 2.3 (Category of strata $\mathcal{S}_{\mathrm{II}}$)**.** We define a category $\mathcal{S}_{\mathrm{II}}$ (depending on data as in Setup II). An object in $\mathcal{S}_{\mathrm{II}}$ is a connected non-empty tree T along with the following labels and decorations. Each edge shall be labeled with a symbol $*(e) \in \{0,1\}$, such that input edges are labeled with 0 and output edges are labeled with 1. Each vertex shall be labeled with a pair of symbols $*^\pm(v) \in \{0,1\}$ such that $*^+(v) \le *^-(v)$ and $*(e^\pm(v)) = *^\pm(v)$ (see Figure 3). There shall also be decorations $\gamma_e \in \mathcal{P}(Y^{*(e)}, \lambda^{*(e)})$ and $\beta_v \in \pi_2(X^{*(v)}, \gamma_{e^+(v)} \sqcup \{\gamma_{e^-(v)}\})$, where $X^{00} = Y^0 := Y^+$, $X^{01} := X$, $X^{11} = Y^1 := Y^-$, as well as basepoints $p_e \in \gamma_e$ for input/output edges. We call v a *symplectization vertex* iff $*^+(v) = *^-(v)$, and we denote by $V_s(T) \subseteq V(T)$ the set of symplectization vertices.

A morphism $T \to T'$ in $\mathcal{S}_{\mathrm{II}}$ consists of a contraction of underlying trees compatible with the decorations, along with paths between the basepoints (see Figures 4 and 5). "Compatible with the decorations" means that $*^+(\pi(v)) \le *^+(v)$, $*^-(\pi(v)) \ge *^-(v)$, and $*(\pi(e)) = *(e)$ for non-contracted edges e along with the conditions from Definition 2.2.

A *concatenation* $\{T_i\}_i$ *in* $\mathcal{S}_{\mathrm{II}}$ shall mean a finite collection of trees $T_i \in \mathcal{S}_{\mathrm{I}}^+ \sqcup \mathcal{S}_{\mathrm{II}} \sqcup \mathcal{S}_{\mathrm{I}}^-$ along with a matching between some pairs of input/output edges of the T_i's (with matching $*(e)$ and γ_e) such that the resulting gluing is a connected tree, along with a choice of paths between the basepoints for each pair of matched edges. Given a concatenation $\{T_i\}_i$ in $\mathcal{S}_{\mathrm{II}}$, there is a resulting object $\#_i T_i \in \mathcal{S}_{\mathrm{II}}$. If $\{T_i\}_i$ is a concatenation in $\mathcal{S}_{\mathrm{II}}$ and $T_i = \#_j T_{ij}$ for

[6] It is helpful to think in terms of the category $\tilde{\mathcal{P}}$ whose objects are Reeb orbits $\gamma \in \mathcal{P}$ together with a basepoint, and whose morphisms are paths between basepoints (as just described). Then the set of isomorphism classes is $|\tilde{\mathcal{P}}| = \mathcal{P}$, and the automorphism group of an object in the isomorphism class of $\gamma \in \mathcal{P}$ is (canonically) $\mathbb{Z}/d_\gamma$.

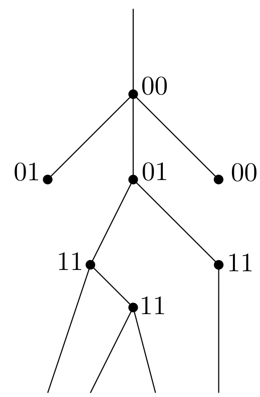

Figure 3: A tree in $\mathcal{S}_{\mathrm{II}}$, with vertex labels as shown. Note that the vertex labels determine the edge labels uniquely (and conversely except for vertices with no outgoing edges).

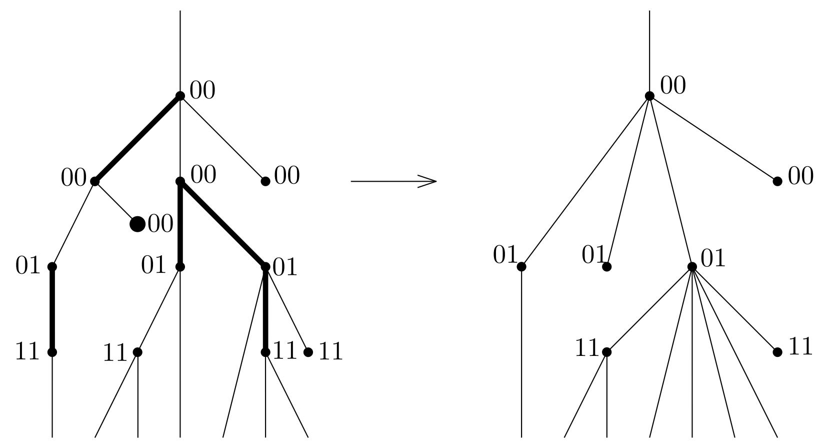

Figure 4: A contraction of trees in $\mathcal{S}_{\mathrm{II}}$. A contraction of trees in $\mathcal{S}_{\mathrm{II}}$ is determined uniquely by the set of contracted edges and the set of vertices whose label changes from 00 to 01 (only needed for vertices with no outgoing edges).

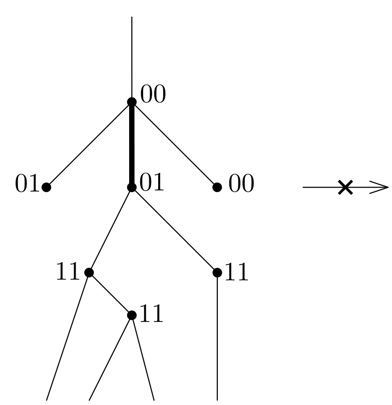

Figure 5: This tree in $\mathcal{S}_{\mathrm{II}}$ cannot be contracted along (exactly) the marked edge (there is no way to consistently label the result).

some concatenations $\{T_{ij}\}_j$ (in whichever of $\mathcal{S}_{\mathrm{II}}$, $\mathcal{S}_{\mathrm{I}}^+$, $\mathcal{S}_{\mathrm{I}}^-$ contains T_i), there is a resulting *composite concatenation* $\{T_{ij}\}_{ij}$ with natural isomorphisms $\#_{ij}T_{ij} = \#_i\#_jT_{ij} = \#_iT_i$.

Call an object $T \in \mathcal{S}_{\mathrm{II}}$ *maximal* iff the only morphism out of T is the identity map (i.e. iff T has exactly one vertex). Every object $T \in \mathcal{S}_{\mathrm{II}}$ has a unique morphism to a maximal object, which is unique up to unqiue isomorphism. Note the following key fact: an object $T \in \mathcal{S}_{\mathrm{II}}$ is maximal iff it cannot be expressed nontrivially as a concatenation.

Definition 2.4 (Category of strata $\mathcal{S}_{\mathrm{III}}$)**.** We define a category $\mathcal{S}_{\mathrm{III}}$ (depending on data as in Setup III). An object in $\mathcal{S}_{\mathrm{III}}$ is a (possibly empty or disconnected) tree T with labels and decorations as in Definition 2.3, along with a set $\mathfrak{s} = \mathfrak{s}(T) \in \{\{0\}, \{1\}, (0,1)\}$.

A morphism $T \to T'$ in $\mathcal{S}_{\mathrm{III}}$ is defined as in Definition 2.3, with the requirement that $\mathfrak{s}(T) \subseteq \overline{\mathfrak{s}(T')}$.

A *concatenation in* $\mathcal{S}_{\mathrm{III}}$ can have the following three types. The first type consists of trees $T_i \in \mathcal{S}_{\mathrm{I}}^+ \sqcup \mathcal{S}_{\mathrm{II}}^{t=0} \sqcup \mathcal{S}_{\mathrm{I}}^-$ with the usual matching data, producing an object $\#_iT_i \in \mathcal{S}_{\mathrm{III}}$ with $\mathfrak{s}(\#_iT_i) = \{0\}$. The second type consists of trees $T_i \in \mathcal{S}_{\mathrm{I}}^+ \sqcup \mathcal{S}_{\mathrm{II}}^{t=1} \sqcup \mathcal{S}_{\mathrm{I}}^-$ with matching data, producing an object $\#_iT_i \in \mathcal{S}_{\mathrm{III}}$ with $\mathfrak{s}(\#_iT_i) = \{1\}$. The third type consists of trees $T_i \in \mathcal{S}_{\mathrm{I}}^+ \sqcup \mathcal{S}_{\mathrm{III}} \sqcup \mathcal{S}_{\mathrm{I}}^-$ (exactly one of which lies in $\mathcal{S}_{\mathrm{III}}$) with matching data, producing an object $\#_iT_i \in \mathcal{S}_{\mathrm{III}}$ (where $\mathfrak{s}(\#_iT_i) = \mathfrak{s}(T_i)$ for the unique $T_i \in \mathcal{S}_{\mathrm{III}}$). A composition of concatenations is defined as before.

Call an object $T \in \mathcal{S}_{\mathrm{III}}$ *maximal* iff the only morphism out of T is the identity map (i.e. iff every component of T has exactly one vertex and $\mathfrak{s}(T) = [0,1]$). Every object $T \in \mathcal{S}_{\mathrm{III}}$ has a unique morphism to a maximal object, which is unique up to unqiue isomorphism. Note the following key fact: an object $T \in \mathcal{S}_{\mathrm{III}}$ is maximal iff it cannot be expressed nontrivially as a concatenation (the "trivial" concatenation means one of the last type, using only $T \in \mathcal{S}_{\mathrm{III}}$).

Definition 2.5 (Category of strata $\mathcal{S}_{\mathrm{IV}}$)**.** We define a category $\mathcal{S}_{\mathrm{IV}}$ (depending on data as in Setup IV). An object in $\mathcal{S}_{\mathrm{IV}}$ is a (possibly empty or disconnected) tree T with the following labels and decorations. Each edge shall be labeled with a symbol $*(e) \in \{0,1,2\}$, such that input edges are labeled with 0 and output edges are labeled with 2. Each vertex shall be labeled with a pair of symbols $*^\pm(v) \in \{0,1,2\}$ such that $*^+(v) \leq *^-(v)$ and $*(e^\pm(v)) = *^\pm(v)$. There shall also be decorations $\gamma_e \in \mathcal{P}(Y^{*(e)})$ and $\beta_v \in \pi_2(X^{*(v)}, \gamma_{e^+(v)} \sqcup \{\gamma_{e^-(v)}\})$ as before. Finally, we specify a set $\mathfrak{s} = \mathfrak{s}(T) \in \{\{0\}, \{\infty\}, (0,\infty)\}$ and we require that if $\mathfrak{s} \in \{\{0\}, (0,\infty)\}$ then $*(v) \in \{00, 02, 22\}$ for all v, and if $\mathfrak{s} = \{\infty\}$ then $*(v) \in \{00, 01, 11, 12, 22\}$ for all v.

A morphism $T \to T'$ in $\mathcal{S}_{\mathrm{IV}}$ is defined as in Definition 2.3, with the requirement that $\mathfrak{s}(T) \subseteq \overline{\mathfrak{s}(T')}$.

A *concatenation in* $\mathcal{S}_{\mathrm{IV}}$ can have the following three types. The first type consists of trees $T_i \in \mathcal{S}_{\mathrm{I}}^0 \sqcup \mathcal{S}_{\mathrm{II}}^{02} \sqcup \mathcal{S}_{\mathrm{I}}^2$ with matching data, producing an object $\#_iT_i \in \mathcal{S}_{\mathrm{IV}}$ with $\mathfrak{s}(\#_iT_i) = \{0\}$. The second type consists of trees $T_i \in \mathcal{S}_{\mathrm{I}}^0 \sqcup \mathcal{S}_{\mathrm{II}}^{01} \sqcup \mathcal{S}_{\mathrm{I}}^1 \sqcup \mathcal{S}_{\mathrm{II}}^{12} \sqcup \mathcal{S}_{\mathrm{I}}^2$ with matching data, producing an object $\#_iT_i \in \mathcal{S}_{\mathrm{IV}}$ with $\mathfrak{s}(\#_iT_i) = \{\infty\}$. The third type consists of trees $T_i \in \mathcal{S}_{\mathrm{I}}^0 \sqcup \mathcal{S}_{\mathrm{IV}} \sqcup \mathcal{S}_{\mathrm{I}}^2$ (exactly one of which lies in $\mathcal{S}_{\mathrm{IV}}$) with matching data, producing an object $\#_iT_i \in \mathcal{S}_{\mathrm{IV}}$ (where $\mathfrak{s}(\#_iT_i) = \mathfrak{s}(T_i)$ for the unique $T_i \in \mathcal{S}_{\mathrm{IV}}$). A composition of concatenations is defined as before.

Call an object $T \in \mathcal{S}_{\mathrm{IV}}$ *maximal* iff the only morphism out of T is the identity map (i.e. iff every component of T has exactly one vertex and $\mathfrak{s}(T) = [0,\infty]$). Every object $T \in \mathcal{S}_{\mathrm{IV}}$ has a unique morphism to a maximal object, which is unique up to unqiue isomorphism. Note the

following key fact: an object $T \in \mathcal{S}_{\mathrm{IV}}$ is maximal iff it cannot be expressed nontrivially as a concatenation (the "trivial" concatenation means one of the last type, using only $T \in \mathcal{S}_{\mathrm{IV}}$).

Definition 2.6 (Slice categories)**.** For $* \in \{\mathrm{I}, \mathrm{II}, \mathrm{III}, \mathrm{IV}\}$, denote by $(\mathcal{S}_*)_{/T}$ the "over-category" whose objects are morphisms $T' \to T$. Similarly define the "under-category" $(\mathcal{S}_*)_{T/}$ (objects are morphisms $T \to T'$) and $(\mathcal{S}_*)_{T//T'}$ (objects are factorizations $T \to T'' \to T'$ of a fixed morphism $T \to T'$). Both $(\mathcal{S}_*)_{T/}$ and $(\mathcal{S}_*)_{T//T'}$ are posets (i.e. have the property that for all (x, y) there is at most one morphism $x \to y$), however $(\mathcal{S}_*)_{/T}$ is essentially never a poset, due to vertices with no outgoing edges (though the set of isomorphism classes $\left|(\mathcal{S}_*)_{/T}\right|$ is a poset).

Definition 2.7 (Automorphism groups)**.** Given a map $T' \to T$, denote by $\operatorname{Aut}(T'/T)$ the subgroup of $\operatorname{Aut}(T')$ compatible with the map $T' \to T$ (i.e. the automorphism group of $(T' \to T) \in (\mathcal{S}_*)_{/T}$). Given a concatenation $\{T_i\}_i$, denote by $\operatorname{Aut}(\{T_i\}_i/\#_i T_i)$ the group of automorphisms of $\{T_i\}_i$ acting trivially on $\#_i T_i$ (i.e. the product $\prod_e \mathbb{Z}/d_{\gamma_e}$ over junction edges, acting diagonally).

2.2 Moduli spaces $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$

We now define the compactified moduli spaces of pseudo-holomorphic curves $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$ relevant for contact homology. We first define the open strata $\mathcal{M}_*(T)$ for each $T \in \mathcal{S}_*$, and then we define the compactified moduli spaces $\overline{\mathcal{M}}_*(T)$ as unions of open strata $\overline{\mathcal{M}}_*(T')$ for $T' \to T$. We also review the definition of the Gromov topology on these moduli spaces, and we recall the fundamental compactness results of [BEHWZ03].

Equip $\mathbb{R} \times S^1$ with coordinates (s, t) and with the standard almost complex structure $j(\partial_s) = \partial_t$, i.e. $z = e^{s+it}$.

Definition 2.8. Fix (Y, λ, J) as in Setup I, and let $u : [0, \infty) \times S^1 \to \hat{Y}$ be a smooth map. We say that u is *positively asymptotic* to a Reeb orbit $\gamma \in \mathcal{P}$ iff:

$$u(s, t) = (Ls + b, \tilde{\gamma}(t)) + o(1) \tag{2.1}$$

as $s \to \infty$, for some $b \in \mathbb{R}$ and some $\tilde{\gamma} : S^1 \to Y$ with $\partial_t \tilde{\gamma} = L \cdot R_\lambda(\tilde{\gamma})$ parameterizing γ. Similarly, we say that $u : (-\infty, 0] \times S^1 \to \hat{Y}$ is *negatively asymptotic* to γ iff it satisfies (2.1) as $s \to -\infty$.

It is straightforward to check that for any map $\phi : [0, \infty) \times S^1 \to [0, \infty) \times S^1$ sending ∞ to ∞ which is a biholomorphism onto its image, u and $u \circ \phi$ have the same asymptotics (by noting that $\phi : D^2 \setminus 0 \to D^2 \setminus 0$ extends holomorphically to $\phi : D^2 \to D^2$). Thus for any Riemann surface C and $p \in C$, it makes sense to say that $u : C \setminus p \to \hat{Y}$ is positively or negative asymptotic to $\gamma \in \mathcal{P}$ at p. Moreover, if $u : C \setminus p \to \hat{Y}$ is asymptotic to $\gamma \in \mathcal{P}$, then it induces a well-defined constant speed parameterization of γ denoted $u_p : S_p C \to Y$, where $S_p C := (T_p C \setminus 0)/\mathbb{R}_{>0}$ is the tangent sphere at p, which is a $U(1)$-torsor due to the complex structure on C.

The notion of being asymptotic to a Reeb orbit generalizes immediately to maps to any $\hat{X}$ from Setups II, III, IV, by virtue of the markings (1.9)–(1.10).

If u is $\hat{J}$-holomorphic, then the property of being positively or negatively asymptotic to a Reeb orbit implies that the error $o(1)$ in (2.1) is $O(e^{-\delta|s|})$ in all derivatives as $|s| \to \infty$, for

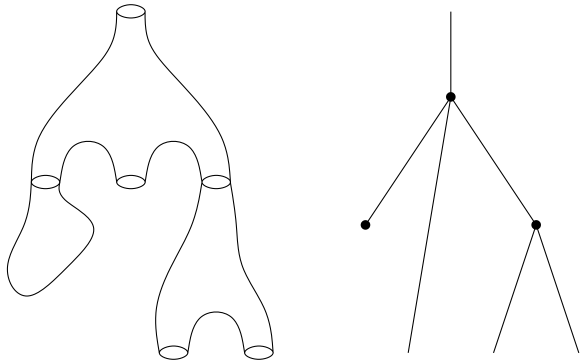

Figure 6: A stable pseudo-holomorphic building and the corresponding tree.

any $\delta < \delta_\gamma$, where $\delta_\gamma > 0$ denotes the least positive eigenvalue of the linearized operator of γ (see Definition 2.16). In particular, it implies the *finite Hofer energy* condition [BEHWZ03, §5.3, §6.1].

Definition 2.9 (Moduli space $\mathcal{M}_{\mathrm{I}}(T)$)**.** A *pseudo-holomorphic building of type* $T \in \mathcal{S}_{\mathrm{I}}$ consists of the following data (see Figure 6):

i. For every vertex v, a closed, connected, nodal Riemann surface of genus zero C_v, along with distinct points $\{p_{v,e} \in C_v\}_e$ indexed by the edges e incident at v.
ii. For every vertex v, a smooth map $u_v : C_v \setminus \{p_{v,e}\}_e \to \hat{Y}$.
iii. We require that u_v be positively asymptotic to $\gamma_{e^+(v)}$ at $p_{v,e^+(v)}$, negatively asymptotic to $\gamma_{e^-(v)}$ at $p_{v,e^-(v)}$, and in the homotopy class β_v.
iv. For every input/output edge e, an "asymptotic marker" $b_e \in S_{p_{v,e}}C$ which is mapped to the basepoint $p_e \in \gamma_e$ by $(u_v)_{p_{v,e}}$.
v. For every interior edge $v \xrightarrow{e} v'$, a "matching isomorphism" $m_e : S_{p_{v,e}}C_v \to S_{p_{v',e}}C_{v'}$ intertwining $(u_v)_{p_{v,e}}$ and $(u_{v'})_{p_{v',e}}$.
vi. We require that u_v be $\hat{J}$-holomorphic, i.e. $(du)^{0,1}_{\hat{J}} = 0$.

An *isomorphism* $(\{C_v\}, \{p_{v,e}\}, \{u_v\}, \{b_e\}, \{m_e\}) \to (\{C'_v\}, \{p'_{v,e}\}, \{u'_v\}, \{b'_e\}, \{m'_e\})$ between pseudo-holomorphic buildings of type T consists of isomorphisms $\{i_v : C_v \to C'_v\}$ and real numbers $\{s_v \in \mathbb{R}\}$ such that $u_v = \tau_{s_v} \circ u'_v \circ i_v$ ($\tau_s : \hat{Y} \to \hat{Y}$ denotes translation by s), $i_v(p_{v,e}) = p'_{v,e}$, $i(b_e) = b'_e$, and $i_{v'} \circ m_e = m'_e \circ i_v$. A pseudo-holomorphic building is called *stable* iff its automorphism group is finite. We denote by $\mathcal{M}_{\mathrm{I}}(T)$ the set of isomorphism classes of stable pseudo-holomorphic buildings of type T.

Definition 2.10 (Moduli space $\mathcal{M}_{\mathrm{II}}(T)$)**.** A *pseudo-holomorphic building of type* $T \in \mathcal{S}_{\mathrm{II}}$ consists of the following data:

i. Same as Definition 2.9(i).
ii. For every vertex v, a smooth map $u_v : C_v \setminus \{p_{v,e}\}_e \to \hat{X}^{*(v)}$.
iii. Same as Definition 2.9(iii).
iv. Same as Definition 2.9(iv).
v. Same as Definition 2.9(v).
vi. Same as Definition 2.9(vi).

An *isomorphism* between pseudo-holomorphic buildings of type T is defined as in Definition 2.9, except that there is a translation $s_v \in \mathbb{R}$ only if v is a symplectization vertex. We denote by $\mathcal{M}_{\mathrm{II}}(T)$ the set of isomorphism classes of stable pseudo-holomorphic buildings of type T.

Definition 2.11 (Moduli space $\mathcal{M}_{\mathrm{III}}(T)$)**.** For $T \in \mathcal{S}_{\mathrm{III}}$, denote by $\mathcal{M}_{\mathrm{III}}(T)$ the union over $t \in \mathfrak{s}(T)$ of the set of isomorphism classes of stable pseudo-holomorphic buildings of type T (as in Definition 2.10) in (X^t, J^t).

Definition 2.12 (Moduli space $\mathcal{M}_{\mathrm{IV}}(T)$)**.** For $T \in \mathcal{S}_{\mathrm{IV}}$, denote by $\mathcal{M}_{\mathrm{IV}}(T)$ the union over $t \in \mathfrak{s}(T)$ of the set of isomorphism classes of stable pseudo-holomorphic buildings of type T (as in Definition 2.10) in (X^{02}_t, J^{02}_t).

Definition 2.13 (Moduli spaces $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$)**.** For $* \in \{\mathrm{I}, \mathrm{II}, \mathrm{III}, \mathrm{IV}\}$, we define:

$$\overline{\mathcal{M}}_*(T) := \bigsqcup_{T' \to T} \mathcal{M}_*(T') / \operatorname{Aut}(T'/T) \tag{2.2}$$

The union is over the poset $|(\mathcal{S}_*)_{/T}|$.

Each of these sets $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$ is equipped with a natural *Gromov topology*. For completeness' sake, let us recall the definition (see also [BEHWZ03, Appendix B]). To define a topology, it suffices to specify a *neighborhood base*[7] at every point. A neighborhood base for the Gromov topology at a given pseudo-holomorphic building (i.e. point in the moduli space) may be described by (1) arbitrarily adding marked points to stabilize the domain, (2) taking a C^0-neighborhood, and (3) forgetting the added marked points. The Gromov topology is Hausdorff, the essential point being that we only consider stable buildings.

Let us now summarize the compactness properties of these spaces $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$ as proved in [BEHWZ03]. For any pseudo-holomorphic building, there is a notion of *Hofer energy*. In broad generality, the moduli space of pseudo-holomorphic buildings with Hofer energy bounded by any fixed constant is compact [BEHWZ03, Theorems 10.1, 10.2, 10.3]. Now the Hofer energy can also be bounded in terms of the homology class and input/output asymptotics of the building [BEHWZ03, Proposition 5.13, 6.3], which lead to the following results. Each space $\overline{\mathcal{M}}_*$ is compact. Moreover, there are only finitely many non-empty moduli spaces $\overline{\mathcal{M}}_*$ for a given fixed γ^+ (or $\{\gamma^+\}$) if the cobordism is exact.

For a morphism $T' \to T$, there is a natural inclusion:

$$\overline{\mathcal{M}}_*(T') / \operatorname{Aut}(T'/T) \hookrightarrow \overline{\mathcal{M}}_*(T) \tag{2.3}$$

so $\overline{\mathcal{M}}_*$ is a functor from $\mathcal{S}_*$ to the category of compact Hausdorff spaces.

For a concatenation $\{T_i\}_i$, there is an induced homeomorphism:

$$\prod_i \overline{\mathcal{M}}_*(T_i) \Big/ \operatorname{Aut}(\{T_i\}_i / \#_i T_i) \to \overline{\mathcal{M}}_*(\#_i T_i) \tag{2.4}$$

[7] A *neighborhood* of a point x in a topological space X is a (not necessarily open) subset $N \subseteq X$ such that $x \in N^\circ$ (the interior of N). A *neighborhood base* at a point $x \in X$ is a collection of neighborhoods $\{N_\alpha\}_\alpha$ of x such that for every open $U \subseteq X$ containing x, there exists some $N_\alpha \subseteq U$. A neighborhood base is necessarily non-empty and *filtered*, i.e. for all α, β, there exists γ with $N_\gamma \subseteq N_\alpha \cap N_\beta$. Conversely, given a set X and for every $x \in X$ a non-empty filtered collection of subsets $\{N^x_\alpha\}_\alpha$ each containing x, there is a unique topology on X such that $\{N^x_\alpha\}_\alpha$ is a neighborhood base at x for all $x \in X$.

Definition 2.14 (Stratification)**.** A *stratification* of a topological space X by a poset $\mathcal{S}$ (usually taken to be finite) is a continuous map $X \to \mathcal{S}$, where $\mathcal{S}$ is endowed with the poset topology, i.e. in which $A \subseteq \mathcal{S}$ is open iff $a \in A \implies \mathcal{S}^{\geq a} \subseteq A$.

Definition 2.15 (Stratifications of moduli spaces)**.** By definition, there is a tautological stratification:

$$\overline{\mathcal{M}}_*(T) \to \left|(\mathcal{S}_*)_{/T}\right| \tag{2.5}$$

As remarked above, each space $\overline{\mathcal{M}}_*(T)$ has only finitely many non-empty strata. The fact that (2.5) is a stratification in the sense of Definition 2.14 follows directly from the definition of the Gromov topology.

2.3 Linearized operators

We now recall the relevant linearized operators associated to the pseudo-holomorphic curves that we consider.

Definition 2.16 (Linearized operator of Reeb orbits)**.** Let $\gamma \in \mathcal{P}$ and choose a parameterization $\tilde{\gamma} : S^1 \to Y$. Consider the (unbounded self-adjoint) operator $D_\gamma : L^2(S^1, \tilde{\gamma}^*\xi) \to L^2(S^1, \tilde{\gamma}^*\xi)$ given by $J\mathcal{L}_{R_\lambda}$. The operator D_γ has only discrete spectrum, and $0 \notin \sigma(D_\gamma)$ is equivalent to the non-degeneracy of γ. Denote by $\delta_\gamma > 0$ the smallest magnitude of an element of $\sigma(D_\gamma)$.

Definition 2.17 (Choices of metric and connection)**.** Let $(\hat{X}, \hat{J})$ be as in Setup I or II. For the purposes of defining function spaces, stating estimates, expressing linearized operators, etc. involving $(\hat{X}, \hat{J})$, we use a Riemannian metric on $\hat{X}$ which is $\mathbb{R}$-invariant in the ends, and we use a $\hat{J}$-linear[8] connection on $T\hat{X}$ which in any end $\hat{Y} \to \hat{X}$ is pulled back from a connection on $T\hat{X}|_Y = TY \oplus \mathbb{R}$. Different choices of metric and connection will always result in uniformly commensurable norms, so the particular choice is not important.

Let C be a compact Riemann surface, and let $\{p_e\}_e$ be a collection of distinct points in C. For the purposes of defining function spaces, stating estimates, etc. involving C, we use a choice of holomorphic cylindrical ends:

$$[0, \infty) \times S^1 \to C \setminus p_e \tag{2.6}$$

near each p_e. We equip C with a Riemannian metric which equals $ds^2 + dt^2$ near p_e, and we equip TC with a j-linear connection for which ∂_s is parallel near p_e. As before, different choices of this data will result in uniformly commensurable norms and estimates, so the particular choice is not important.

Definition 2.18 (Weighted Sobolev spaces $W^{k,2,\delta}$)**.** We recall the weighted Sobolev spaces:

$$W^{k,2,\delta}(C, u^*T\hat{X}) \tag{2.7}$$

$$W^{k,2,\delta}(C, u^*T\hat{X}_{\hat{J}} \otimes_{\mathbb{C}} \Omega^{0,1}_C) \tag{2.8}$$

[8]Meaning $\hat{J}\nabla_X Y = \nabla_X \hat{J}Y$, i.e. $\nabla\hat{J} = 0$.

where $u : C \setminus \{p_e\}_e \to \hat{X}$ is a smooth map which in a neighborhood of each puncture p_e is $\hat{J}$-holomorphic and asymptotic to a Reeb orbit $\gamma_e \in \mathcal{P}^{\pm}$. These spaces are defined for integers $k \geq 0$ and *admissible* real numbers δ, meaning that $\delta < \min(\delta_\gamma, 1)$ (minimum over all asymptotic orbits γ_e).

The $W^{k,2,\delta}$-norm is defined as follows. Away from $\{p_e\}_e$, we use the usual $W^{k,2}$-norm, and near a given p_e, the contribution to the norm squared is given by:

$$\int_{[0,\infty)\times S^1} \sum_{j=0}^{k} \left|D^j f\right|^2 e^{2\delta s}\, ds\, dt \tag{2.9}$$

Equivalently up to uniform commensurability, one can set $\|f\|_{k,2,\delta} := \|\mu \cdot f\|_{k,2}$ for some smooth function μ which equals 1 away from the ends and which equals $e^{\delta|s|}$ in any end. The $W^{k,2,\delta}$-norm is independent up to uniform commensurability of the choices involved in its definition, due to the decay of (2.1) and admissibility of δ.

Definition 2.19. Let $\overline{\mathcal{M}}_{g,n+m_{(2)}}$ denote the moduli space of compact genus g nodal Riemann surfaces with $n+m$ marked points and trivializations $\mathbb{C} \xrightarrow{\sim} T_{p_i}C$ at the last m marked points. Clearly there is a natural forgetful map $\overline{\mathcal{M}}_{g,n+m_{(2)}} \to \overline{\mathcal{M}}_{g,n+m}$ which is a principal $\mathrm{GL}_1(\mathbb{C})^m$-bundle. In particular, $\overline{\mathcal{M}}_{g,n+m_{(2)}}$ is a complex orbifold (and a complex manifold for $g = 0$).

Definition 2.20 (Weighted Sobolev spaces $\tilde{W}^{k,2,\delta}$)**.** We define a weighted Sobolev space:

$$\tilde{W}^{k,2,\delta}(C, u^*T\hat{X}) \tag{2.10}$$

Roughly speaking, $\tilde{W}^{k,2,\delta}$ equals $W^{k,2,\delta}$ plus the space of variations in the almost complex structure on C.

If $2\#\{p_e\} - 3 \leq 0$, then we set:

$$\tilde{W}^{k,2,\delta}(C, u^*T\hat{X}) := W^{k,2,\delta}(C, u^*T\hat{X}) \Big/ \ker\Big(\mathfrak{aut}(C, \{p_e\}_e) \to \bigoplus_e \mathfrak{gl}(T_{p_e}C)\Big) \tag{2.11}$$

Here $\ker\big(\mathfrak{aut}(C, \{p_e\}_e) \to \bigoplus_e \mathfrak{gl}(T_{p_e}C)\big)$ denotes the Lie algebra of the group of automorphisms of C which fix each p_e and act as the identity on each $T_{p_e}C$; this Lie algebra of vector fields on C maps into $W^{k,2,\delta}(C, u^*T\hat{X})$ by differentiating u (since δ is admissible).

If $2\#\{p_e\} - 3 \geq 0$, then we set:

$$\tilde{W}^{k,2,\delta}(C, u^*T\hat{X}) := W^{k,2,\delta}(C, u^*T\hat{X}) \oplus V \tag{2.12}$$

where $V \subseteq C_c^\infty(C \setminus \{p_e\}_e, \mathrm{End}^{0,1}(TC))$ (i.e. the space of infinitesimal deformations of the almost complex structure on C supported away from $\{p_e\}_e$) is a subspace projecting isomorphically onto the tangent space to $\overline{\mathcal{M}}_{0,0+(\{p_e\}_e)_{(2)}}$ at C. Every other such subspace V' may be obtained as the image of $v \mapsto v + \mathcal{L}_{X(v)}j$ for some unique $X : V \to C^\infty(C, TC)$ with $X = 0$ and $dX = 0$ at p_e, and the two resulting spaces (2.12) are canonically isomorphic via $(\xi, v) \mapsto (\xi + X(v)u, v + \mathcal{L}_{X(v)}j)$.

Definition 2.21 (Linearized operators)**.** Let $u : C \to \hat{X}$ be as in Definition 2.18. There is a linearized operator:

$$W^{k,2,\delta}(C \setminus \{p_e\}, u^*TX) \oplus V \to W^{k-1,2,\delta}(\tilde{C} \setminus \{p_e\}, u^*T\hat{X}_{\hat{j}} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}}) \tag{2.13}$$

expressing the first order change in $(du)^{0,1}$ as u and j vary. Here $V \subseteq C^\infty_c(C\setminus\{p_e\}, \operatorname{End}^{0,1}(TC))$ is as in Definition 2.20, and $\tilde{C}$ denotes the normalization of C, i.e. the unique smooth compact Riemann surface equipped with a map $\tilde{C} \to C$ which identifies points in pairs to form the nodes of C. This linearized operator is Fredholm (see, e.g. Lockhart–McOwen [LM85]).

The operator (2.13) depends on a choice of $\hat{J}$-linear connection as in Definition 2.17. If u is everywhere $\hat{J}$-holomorphic, then the linearized operator is independent of the choice of $\hat{J}$-linear connection and descends to $\tilde{W}^{k,2,\delta}$.

Definition 2.22 (Linearized operators)**.** Given a pseudo-holomorphic building of type T, we have defined a linearized operator:

$$\bigoplus_{v\in V(T)} \tilde{W}^{k,2,\delta}(C_v, u_v^*T\hat{X}_v) \to \bigoplus_{v\in V(T)} W^{k-1,2,\delta}(\tilde{C}_v, u_v^*(T\hat{X}_v)_{\hat{J}_v} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}_v}) \tag{2.14}$$

A point in a moduli space $\overline{\mathcal{M}}_*$ is called *regular* iff the linearized operator (2.14) of the corresponding pseudo-holomorphic building is surjective (it follows from elliptic regularity theory that this condition is independent of k and δ). A moduli space $\overline{\mathcal{M}}_*$ is called regular iff all of its points are regular.

Remark 2.23. For $* \in \{\mathrm{III}, \mathrm{IV}\}$, there is a less restrictive notion of regularity, where we also allow variations in $t \in \mathfrak{s}(T)$ (note that this does *not* include variations in t if $\mathfrak{s}(T) = \{0\}$, for example). In the statements of Theorems III, IV, we mean regularity in this less restrictive sense. However, for the rest of the paper, regularity shall mean regularity in the sense of Definition 2.22 above.

Lemma 2.24. *Fix a Reeb orbit γ and consider the trivial cylinder $(\mathrm{id}\times\gamma) : C = \mathbb{R}\times S^1 \to \mathbb{R}\times Y$. The associated linearized operator is an isomorphism (and in particular is surjective).*

Proof. The linearized operator D may be decomposed into the tangential and normal deformation operators D_T and D_N. Precisely, there is a diagram whose rows are short exact sequences:

$$\begin{tikzcd}
\tilde{W}^{k,2,\delta}(C, TC) \arrow[r] \arrow[d, "D_T"] & \tilde{W}^{k,2,\delta}(C, \gamma^*TY \oplus \mathbb{R}\partial_s) \arrow[r] \arrow[d, "D"] & W^{k,2,\delta}(C, \gamma^*\xi) \arrow[d, "D_N"] \\
W^{k-1,2,\delta}(C, TC \otimes_{\mathbb{C}} \Omega^{0,1}_C) \arrow[r] & W^{k-1,2,\delta}(C, (\gamma^*TY \oplus \mathbb{R}\partial_s) \otimes_{\mathbb{C}} \Omega^{0,1}_C) \arrow[r] & W^{k-1,2,\delta}(C, \gamma^*\xi \otimes_{\mathbb{C}} \Omega^{0,1}_C)
\end{tikzcd} \tag{2.15}$$

Note that the domain of D_T includes variations in complex structure on C, while the domain of D_N does not. It suffices to show that both D_T and D_N are isomorphisms.

The complex $W^{k,2,\delta}(C, TC) \to W^{k-1,2,\delta}(C, TC \otimes_{\mathbb{C}} \Omega^{0,1}_C)$ calculates $H^\bullet(\mathbb{P}^1, T_{\mathbb{P}^1}(-2[0] - 2[\infty]))$, which is concentrated in degree one. By definition $\tilde{W}^{k,2,\delta}$ is $W^{k,2,\delta}$ direct sum a space which maps isomorphically to $H^1(\mathbb{P}^1, T_{\mathbb{P}^1}(-2[0] - 2[\infty])) = T_{\mathbb{P}^1,\{0,\infty\}}\mathcal{M}_{0,0+2_{(2)}}$. Thus D_T is an isomorphism.

The operator D_N may be expressed as $\partial_s + J\nabla_t$ where $\nabla_t := \mathcal{L}_{R_\lambda}$ is the connection on $\gamma^*\xi$. It is an isomorphism since δ is admissible. ∎

2.4 Index of moduli spaces

We now define a notion of *index* and *codimension* for objects of $\mathcal{S}_{\mathrm{I}}, \mathcal{S}_{\mathrm{II}}, \mathcal{S}_{\mathrm{III}}, \mathcal{S}_{\mathrm{IV}}$.

Definition 2.25 (Index $\mu(T)$)**.** We define $\mu(T)$ to be the Fredholm index of (2.14) (for non-nodal C_v). Note that (2.14) makes sense for any $\{u_v\}_v$ which approach trivial cylinders sufficiently rapidly and varies nicely in families of such $\{u_v\}_v$, hence $\mu(T)$ is well-defined.

Standard arguments allow one to express $\mu(T)$ in terms of the Conley–Zehnder indices of the Reeb orbits γ_{e^+} and $\{\gamma_{e^-}\}$ and the homology class of β (see [EGH00, Proposition 1.7.1] or [BM04, Proposition 4]):

$$\mu(T) = [\mathrm{CZ}(\gamma^+) + n - 3] - \sum_i [\mathrm{CZ}(\gamma_i^-) + n - 3] \tag{2.16}$$

(where CZ is the Conley–Zehnder index relative to the homology class of β). We thus define $|\gamma| := \mathrm{CZ}(\gamma) + n - 3 \in \mathbb{Z}/2c_1(\xi) \cdot H_2(Y)$.

The index satisfies $\mu(T) = \mu(T')$ for any morphism $T \to T'$ (this is evident from the formula in terms of Conley–Zehnder indices), and is additive under concatenations, that is $\mu(\#_i T_i) = \sum_i \mu(T_i)$ (trivial by definition).

Definition 2.26 (Codimension $\operatorname{codim}(T)$)**.** We define:

$$\operatorname{codim}(T) := \#V_s(T) - \dim \mathfrak{s}(T) \tag{2.17}$$

(recall that $V_s(T) \subseteq V(T)$ denotes the symplectization vertices, i.e. those v for which $*^+(v) = *^-(v)$).

For a morphism $T \to T'$, let $\operatorname{codim}(T/T') := \operatorname{codim} T - \operatorname{codim} T'$. Note that $\operatorname{codim}(T/T') \geq 0$, with equality iff the map is an isomorphism. Also note that $\operatorname{codim}(T/T') > 1$ iff the map $T \to T'$ factors nontrivially.

Definition 2.27 (Virtual dimension $\operatorname{vdim}(T)$)**.** We define:

$$\begin{aligned} \operatorname{vdim}(T) &:= \mu(T) - \#V_s(T) + \dim \mathfrak{s}(T) \\ &= \mu(T) - \operatorname{codim} T \end{aligned} \tag{2.18}$$

(this is the "expected dimension" of $\overline{\mathcal{M}}_*(T)$).

2.5 Orientations of moduli spaces

We now review the theory of orientations in contact homology. The general analytic methods used to orient moduli spaces of pseudo-holomorphic curves were introduced by Floer–Hofer [FH93] (see also Bourgeois–Mohnke [BM04]). The resulting algebraic structure relevant for contact homology was worked out by Eliashberg–Givental–Hofer [EGH00] (see also Bourgeois–Mohnke [BM04]).

For every Reeb orbit γ with basepoint p (resp. object $T \in \mathcal{S}_{\mathrm{I}}, \mathcal{S}_{\mathrm{II}}, \mathcal{S}_{\mathrm{III}}, \mathcal{S}_{\mathrm{IV}}$), we will define an *orientation line* $\mathfrak{o}_\gamma$ (resp. $\mathfrak{o}_T^0$). We shall see that the virtual orientation sheaf of $\overline{\mathcal{M}}_*(T)$ is canonically isomorphic to:

$$\mathfrak{o}_T := \mathfrak{o}_T^0 \otimes (\mathfrak{o}_{\mathbb{R}}^\vee)^{\otimes V_s(T)} \otimes \mathfrak{o}_{\mathfrak{s}(T)} \tag{2.19}$$

Recall that an *orientation line* means a $\mathbb{Z}/2$-graded free $\mathbb{Z}$-module of rank one.

Definition 2.28. For a Fredholm map $A : E \to F$, we denote by $[A]$ the virtual vector space $\ker A - \operatorname{coker} A$. By $\mathfrak{o}(V) = \mathfrak{o}_V$ we mean the orientation line[9] of the vector space V, i.e. $H_{\dim V}(V, V \setminus 0)$, and for a virtual vector space we set $\mathfrak{o}(V - V') := \mathfrak{o}_V \otimes \mathfrak{o}_{V'}^\vee$.

Definition 2.29 (Orientation lines $\mathfrak{o}_{\gamma,p}$ of Reeb orbits)**.** Let $\gamma \in \mathcal{P} = \mathcal{P}(Y, \lambda)$, and fix a constant speed parameterization $\tilde{\gamma} : S^1 \to Y$ of γ (equivalently, fix a basepoint $p = \tilde{\gamma}(0) \in \gamma$). We consider the bundle $V := \tilde{\gamma}^*\xi \oplus \mathbb{C}$ over $[0, \infty) \times S^1 \subseteq \mathbb{C}$. The bundle V is equipped with a connection, namely the connection on $\tilde{\gamma}^*\xi$ induced by the Lie derivative $\mathcal{L}_{R_\lambda}$ plus the trivial connection on $\mathbb{C}$. Now extend the pair $(V, \bar{\partial})$ to all of $\mathbb{C}$, and define:

$$\mathfrak{o}_{\gamma,p} := \mathfrak{o}([W^{k,2,\delta}(\mathbb{C}, V) \to W^{k-1,2,\delta}(\mathbb{C}, V \otimes_{\mathbb{C}} \Omega^{0,1}_{\mathbb{C}})]) \tag{2.20}$$

Now $\mathfrak{o}_{\gamma,p}$ is independent of the choice of extension of $(V, \bar{\partial})$ up to unique isomorphism, as can be seen as follows. The space of extensions of $(V, \bar{\partial})$ is homotopy equivalent to $\operatorname{Maps}(S^2, BU(n))$ (noting that $\pi_1(BU(n)) = \pi_0(U(n)) = 0$ so V is trivial over S^1) and $\pi_i \operatorname{Maps}(S^2, BU(n)) = \pi_{i+2}(BU(n)) = \pi_{i+1}(U(n))$; in particular $\pi_0 = \mathbb{Z}$ and $\pi_1 = 0$. By simple connectivity, the line $\mathfrak{o}_{\gamma,p}$ depends at most on the choice of connected component of $\operatorname{Maps}(S^2, BU(n))$ (classified by relative Chern class). To identify (canonically) the $\mathfrak{o}_{\gamma,p}$ from different Chern classes, argue as follows. Fix $z \in \mathbb{C}$ and let $V(-z)$ denote the complex vector bundle whose smooth sections are those smooth sections f of V satisfying $f(z) = 0$ and $\bar{\partial} f(z) = 0$ (i.e. whose $(0,1)$-jet at z vanishes). Now $V(-z)$ naturally inherits the $\bar{\partial}$-operator from V, and there is a diagram with exact rows:

$$\begin{array}{ccccccccc} 0 & \longrightarrow & W^{k,2,\delta}(\mathbb{C}, V(-z)) & \longrightarrow & W^{k,2,\delta}(\mathbb{C}, V) & \longrightarrow & J^{0,1}_z(V) & \longrightarrow & 0 \\ & & \downarrow & & \downarrow & & \downarrow & & \\ 0 & \longrightarrow & W^{k-1,2,\delta}(\mathbb{C}, V(-z) \otimes_{\mathbb{C}} \Omega^{0,1}_{\mathbb{C}}) & \longrightarrow & W^{k-1,2,\delta}(\mathbb{C}, V \otimes_{\mathbb{C}} \Omega^{0,1}_{\mathbb{C}}) & \longrightarrow & V_z \otimes \Omega^{0,1}_{\mathbb{C},z} & \longrightarrow & 0 \end{array} \tag{2.21}$$

Applying the snake lemma yields the desired identification, since the rightmost vertical map is surjective with kernel V_z which as a complex vector space is canonically oriented (and the resulting identification is independent of $z \in \mathbb{C}$ by a homotopy argument).

Any rotation $\phi : S^1 \to S^1$ with $\tilde{\gamma} \circ \phi = \tilde{\gamma}'$ (equivalently, any path between basepoints $p \to p'$ in the sense of Definition 2.2) gives rise to an isomorphism $\mathfrak{o}_{\gamma,p} \to \mathfrak{o}_{\gamma,p'}$ (by functoriality of the construction of $\mathfrak{o}_{\gamma,p}$).

Definition 2.30 (Parity of Reeb orbits)**.** The *parity* $|\gamma| \in \mathbb{Z}/2$ of $\gamma \in \mathcal{P}$ is the parity of $\mathfrak{o}_\gamma$. By (2.16) we have $|\gamma| = \mathrm{CZ}(\gamma) + n - 3$; this can also be expressed as $|\gamma| = \operatorname{sign}(\det(I - A_\gamma)) \in \{\pm 1\} = \mathbb{Z}/2$ where A_γ denotes the linearized return map of γ acting on ξ_p (using the general property of the Conley–Zehnder index $(-1)^{\mathrm{CZ}(\Psi_t)} = (-1)^{n-1} \operatorname{sign}(\det(I - \Psi_1))$ for $\Psi_t \in \mathrm{Sp}_{2n-2}(\mathbb{R})$). It thus follows that:

$$|\gamma| = \#(\sigma(A_\gamma) \cap (0, 1)) \in \mathbb{Z}/2 \tag{2.22}$$

[9]One should be careful to distinguish the orientation line $\mathfrak{o}_V$ from the determinant line $\det V := \wedge^{\dim V} V$. There is no functorial isomorphism $\mathfrak{o}_V \otimes_{\mathbb{Z}} \mathbb{R} = \det V$, though of course $\mathfrak{o}_V = \mathfrak{o}_{\det V}$.

where $\sigma(\cdot)$ denotes the spectrum (recall that the spectrum of any matrix $A \in \mathrm{Sp}_{2n}(\mathbb{R})$ lies in $\mathbb{R}^\times \cup \{z \in \mathbb{C} : |z| = 1\}$ and that $1 \notin \sigma(A_\gamma)$ is equivalent to the non-degeneracy of γ). Note that by definition γ is non-degenerate iff $1 \notin \sigma(A_\gamma)$. It follows that the index of the k-fold multiple cover γ_k of γ is given by:

$$|\gamma_k| = |\gamma| + (k+1)\#(\sigma(A_\gamma) \cap (-1, 0)) \in \mathbb{Z}/2 \tag{2.23}$$

Definition 2.31 (Good and bad Reeb orbits)**.** There is an action of $\mathbb{Z}/d_\gamma$ on $\mathfrak{o}_{\gamma,p}$ by functoriality, which just amounts to a homomorphism $\mathbb{Z}/d_\gamma \to \{\pm 1\}$ (independent of p since $\mathbb{Z}/d_\gamma$ is abelian). The orbit γ is called *good* iff this homomorphism is trivial (and *bad* otherwise). For good γ, we thus have an orientation line $\mathfrak{o}_\gamma$ independent of p up to unique isomorphism.

The bad Reeb orbits are precisely the even multiple covers γ_{2k} of *simple* orbits γ with $\#(\sigma(A_\gamma) \cap (-1, 0))$ odd by [EGH00, Lemma 1.8.8, Remark 1.9.2]. To see this, it suffices to show that a generator of the $\mathbb{Z}/k$ action on $\mathfrak{o}_{\gamma_k}$ acts by $|\gamma_k| - |\gamma|$. This can be proven by pulling back the operator from (2.20) under $z \mapsto z^k$ and analyzing the representations of $\mathbb{Z}/k$ occurring in the kernel and cokernel (see [BM04, Proof of Theorem 3]).

Definition 2.32 (Orientation lines $\mathfrak{o}_T^0$ of trees)**.** For any $T \in \mathcal{S}_{\mathrm{I}}, \mathcal{S}_{\mathrm{II}}, \mathcal{S}_{\mathrm{III}}, \mathcal{S}_{\mathrm{IV}}$, we define the orientation line $\mathfrak{o}_T^0$ to be the orientation line of (2.14) (for non-nodal C_v). Note that (2.14) makes sense for any $\{u_v\}_v$ which approach trivial cylinders sufficiently rapidly, and varies nicely in families of such $\{u_v\}_v$. Note also that the space of "abstract" Cauchy–Riemann operators with fixed asymptotics and fixed relative Chern class is simply connected since $\pi_2(U(n)) = 0$ (as in Definition 2.29), so $\mathfrak{o}_T^0$ is well-defined. Note also that we may use $W^{k,2,\delta}$ in place of $\tilde{W}^{k,2,\delta}$ in this definition, since their "difference" is a complex vector space and thus is canonically oriented.

For any morphism $T \to T'$, there is a canonical isomorphism $\mathfrak{o}_T^0 \to \mathfrak{o}_{T'}^0$, defined by the "kernel gluing" operation introduced in Floer–Hofer [FH93] (see also Bourgeois–Mohnke [BM04]). This makes $\mathfrak{o}^0$ into a functor from $\mathcal{S}_{\mathrm{I}}$ to the category of orientation lines and isomorphisms. Note that gluing the linearized operators with domain $W^{k,2,\delta}$ results in an index increase of two, gluing the linearized operators with domain $\tilde{W}^{k,2,\delta}$ preserves the index, and gluing the linearized operators with domain $W^{k,2,\delta}$ plus decay to constants in the ends results in an index decrease of two (the kernel gluing operation can be done in any of these contexts; Bourgeois–Mohnke choose the last).

For any concatenation $\{T_i\}$, there is a tautological identification $\mathfrak{o}_{\#_i T_i}^0 = \bigotimes_i \mathfrak{o}_{T_i}^0$.

The Floer–Hofer kernel gluing operation also lets us glue the operators (2.20) for $\{\gamma_{e^-}\}$ to the operator for T, producing an operator (2.20) for γ_{e^+}. This gives a canonical isomorphism:

$$\mathfrak{o}_T^0 = \mathfrak{o}_{\gamma_{e^+}, p_{e^+}} \otimes \mathfrak{o}_{\Gamma_{e^-}, p_{e^-}}^\vee \tag{2.24}$$

Moreover, these identifications are compatible with the identifications $\mathfrak{o}_{\#_i T_i}^0 = \bigotimes_i \mathfrak{o}_{T_i}^0$ due to the associativity of the gluing map. Note that this argument relies on the fact that the topology of the curves in question is particularly simple; to prove the analogous result in the SFT setting requires another step (observing that one can make the operator $\mathbb{C}$-linear in ends; see Bourgeois–Mohnke [BM04, Proposition 8]). The parity of $\mathfrak{o}_T^0$ equals $\mu(T)$ by definition.

For any $T \in \mathcal{S}_{\mathrm{I}}, \mathcal{S}_{\mathrm{II}}, \mathcal{S}_{\mathrm{III}}, \mathcal{S}_{\mathrm{IV}}$, the orientation line:

$$(\mathfrak{o}_{\mathbb{R}}^{\vee})^{\otimes V_s(T)} \otimes \mathfrak{o}_{\mathfrak{s}(T)} \tag{2.25}$$

plays an important role (here $\mathfrak{o}_{\mathfrak{s}(T)}$ denotes the global sections of the orientation sheaf of $\mathfrak{s}$ considered as a real manifold). Its parity is clearly $-\operatorname{codim} T$. For any morphism $T' \to T$ of codimension one, there is a canonical odd "boundary" map:

$$(\mathfrak{o}_{\mathbb{R}}^{\vee})^{\otimes V_s(T)} \otimes \mathfrak{o}_{\mathfrak{s}(T)} \to (\mathfrak{o}_{\mathbb{R}}^{\vee})^{\otimes V_s(T')} \otimes \mathfrak{o}_{\mathfrak{s}(T')} \tag{2.26}$$

For any concatenation $\{T_i\}$, there is a tautological identification:

$$(\mathfrak{o}_{\mathbb{R}}^{\vee})^{\otimes V_s(\#_i T_i)} \otimes \mathfrak{o}_{\mathfrak{s}(\#_i T_i)} = \bigotimes_i (\mathfrak{o}_{\mathbb{R}}^{\vee})^{\otimes V_s(T_i)} \otimes \mathfrak{o}_{\mathfrak{s}(T_i)} \tag{2.27}$$

We shall see that the virtual orientation sheaf of $\overline{\mathcal{M}}_*(T)$ is canonically isomorphic to:

$$\mathfrak{o}_T := \mathfrak{o}_T^0 \otimes (\mathfrak{o}_{\mathbb{R}}^{\vee})^{\otimes V_s(T)} \otimes \mathfrak{o}_{\mathfrak{s}(T)} \tag{2.28}$$

Note that the parity of $\mathfrak{o}_T$ equals $\operatorname{vdim}(T)$.

2.6 Local models G_{I}, G_{II}, G_{III}, G_{IV}

We now define spaces G_{I}, G_{II}, G_{III}, G_{IV} which will serve as local models for the regular loci in the moduli spaces $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$ (and their "thickenings", to be introduced in §3). The space $(G_*)_{T/}$ for $T \in \mathcal{S}_*$ should be thought of as the space of possible gluing parameters for a pseudo-holomorphic building of type T.

Definition 2.33 (Cell-like stratification [Par15, Definition 6.1.2])**.** Let $(X, \partial X)$ be a topological manifold with boundary with stratification by $(\mathcal{S}, \partial\mathcal{S})$, along with an order preserving function $\dim : \mathcal{S} \to \mathbb{Z}$. We say this stratification is *cell-like* iff $(X^{\leq \mathfrak{s}}, X^{<\mathfrak{s}})$ is a topological manifold with boundary of dimension $\dim \mathfrak{s}$ for all $\mathfrak{s} \in \mathcal{S}$.

For stratifications by the poset $(\mathcal{S}_*)_{T/}$, we take the dimension function to be:

$$(T \to T') \mapsto \#V_s(T) - \#V_s(T') + \dim \mathfrak{s}(T') = \operatorname{codim}(T/T') + \dim \mathfrak{s}(T) \tag{2.29}$$

Remark 2.34. The spaces G_* considered here are "manifolds with generalized corners" as introduced by Joyce [Joy15a] (see also Kottke–Melrose [KM15] and Gillam–Molcho [GM15]). This is most apparent after the change of variables $h = e^{-g}$.

Definition 2.35 (Space G_{I})**.** Let $T \in \mathcal{S}_{\mathrm{I}}$. We define:

$$(G_{\mathrm{I}})_{T/} := (0, \infty]^{E^{\mathrm{int}}(T)} \tag{2.30}$$

There is a natural stratification $(G_{\mathrm{I}})_{T/} \to (\mathcal{S}_{\mathrm{I}})_{T/}$, sending $g = \{g_e\}_e$ to the map $T \to T'$ which contracts those edges $e \in E^{\mathrm{int}}(T)$ for which $g_e < \infty$.

Denote by $0 \in (G_{\mathrm{I}})_{T/}$ the point corresponding to all gluing parameters equal to ∞ (i.e. corresponding to no gluing at all).

Lemma 2.36. *The stratification $(G_{\mathrm{I}})_{T/} \to (\mathcal{S}_{\mathrm{I}})_{T/}$ is cell-like.*

Proof. By inspection. $\square$

Definition 2.37 (Space G_{II})**.** Let $T \in \mathcal{S}_{\mathrm{II}}$. We define:

$$(G_{\mathrm{II}})_{T/} := \Big\{(\{g_e\}_e, \{g_v\}_v) \in (0,\infty]^{E^{\mathrm{int}}(T)} \times (0,\infty]^{V_{00}(T)} \;\Big|\; g_v = g_e + g_{v'} \text{ for } v \xrightarrow{e} v' \text{ and } v \in V_{00}(T)\Big\} \tag{2.31}$$

We interpret $g_{v'} = 0$ if $v' \notin V_{00}(T)$. Here $V_{ij}(T) \subseteq V(T)$ is the subset of v with $(*^+(v), *^-(v)) = (i,j)$.

There is a natural stratification $(G_{\mathrm{II}})_{T/} \to (\mathcal{S}_{\mathrm{II}})_{T/}$, sending $g = (\{g_e\}_e, \{g_v\}_v)$ to the map $T \to T'$ which contracts those edges e for which $g_e < \infty$, and changes $*(v) = 00$ to $*(v) = 01$ for those vertices v with $g_v < \infty$.

Denote by $0 \in (G_{\mathrm{II}})_{T/}$ the point corresponding to all gluing parameters equal to ∞ (i.e. corresponding to no gluing at all).

Lemma 2.38. *The stratification $(G_{\mathrm{II}})_{T/} \to (\mathcal{S}_{\mathrm{II}})_{T/}$ is cell-like.*

Proof. Denote by $(G_{\mathrm{II}})_{T//T'}$ the inverse image of $(\mathcal{S}_{\mathrm{II}})_{T//T'} \subseteq (\mathcal{S}_{\mathrm{II}})_{T/}$. Now observe that for any $f : T \to T'$, we have:

$$(G_{\mathrm{II}})_{T//T'} = \prod_{*(v')=00} (G_{\mathrm{I}}^+)_{f^{-1}(v')/} \times \prod_{*(v')=01} (G_{\mathrm{II}})_{f^{-1}(v')/} \times \prod_{*(v')=11} (G_{\mathrm{I}}^-)_{f^{-1}(v')/} \tag{2.32}$$

$$(\mathcal{S}_{\mathrm{II}})_{T//T'} = \prod_{*(v')=00} (\mathcal{S}_{\mathrm{I}}^+)_{f^{-1}(v')/} \times \prod_{*(v')=01} (\mathcal{S}_{\mathrm{II}})_{f^{-1}(v')/} \times \prod_{*(v')=11} (\mathcal{S}_{\mathrm{I}}^-)_{f^{-1}(v')/} \tag{2.33}$$

compatibly with stratifications. Note that a product of cell-like stratifications is cell-like. Thus by induction (say, on the number of vertices of T), it suffices to show that $(G_{\mathrm{II}})_{T/}$ is a topological manifold with boundary, whose interior coincides with the top stratum.

To show that $(G_{\mathrm{II}})_{T/}$ is a topological manifold with boundary, we do a change of variables $h = e^{-g} \in [0,1)$. For convenience, we will allow $h \in [0,\infty)$ (this relaxation is certainly permitted for the present purpose). The relation $g_v = g_e + g_{v'}$ now becomes $h_v = h_e h_{v'}$. Under this relation $h_v = h_e h_{v'}$, observe that $h_v \in [0,\infty)$ and $h_e^2 - h_{v'}^2 \in (-\infty,\infty)$ determine $h_e \in [0,\infty)$ and $h_{v'} \in [0,\infty)$ uniquely, since:

$$(h_e + i h_{v'})^2 = (h_e^2 - h_{v'}^2) + 2i h_v \tag{2.34}$$

Thus if we perform another change of variables $q_e = h_e^2 - h_{v'}^2$ for $v \xrightarrow{e} v'$, then we have:

$$(G_{\mathrm{II}})_{T/} = \left\{\begin{matrix} h_{v^{\mathrm{top}}} \in [0,\infty) & \text{if } *(v^{\mathrm{top}}) = 00 \\ q_e \in (-\infty,\infty) & \text{for } v \xrightarrow{e} v' \text{ with } *(v') = 00 \\ h_e \in [0,\infty) & \text{for } *(e) = 1 \end{matrix}\right\} \tag{2.35}$$

This is clearly a topological manifold with boundary of dimension $\operatorname{codim} T$. Furthermore, the top stratum is the locus where $h_{v^{\mathrm{top}}} > 0$ and $h_e > 0$, which is clearly its interior.

Though not logically necessary due to the inductive reasoning above, let us remark that one can also express the stratification of $(G_{\mathrm{II}})_{T/}$ by $(\mathcal{S}_{\mathrm{II}})_{T/}$ concretely in terms of the coordinates (2.35) and thereby verify explicitly that it is cell-like. $\square$

Definition 2.39 (Space G_{III})**.** Let $T \in \mathcal{S}_{\mathrm{III}}$. Let:

$$(G_{\mathrm{III}})_{T/} := \Big\{(\{g_e\}_e, \{g_v\}_v) \in (0,\infty]^{E^{\mathrm{int}}(T)} \times (0,\infty]^{V_{00}(T)} \;\Big|\; g_v = g_e + g_{v'} \text{ for } v \xrightarrow{e} v' \text{ and } v \in V_{00}(T)\Big\} \times \begin{cases} [0,1) & \mathfrak{s}(T) = \{0\} \\ (0,1) & \mathfrak{s}(T) = (0,1) \\ (0,1] & \mathfrak{s}(T) = \{1\} \end{cases} \tag{2.36}$$

(the first factor is identical to (2.31)).

There is a natural stratification $(G_{\mathrm{III}})_{T/} \to (\mathcal{S}_{\mathrm{III}})_{T/}$, sending $g = (\{g_e\}_e, \{g_v\}_v, t)$ to the map $T \to T'$ which contracts those edges e for which $g_e < \infty$, changes $*(v) = 00$ to $*(v) = 01$ for those vertices v with $g_v < \infty$, and has $t \in \mathfrak{s}(T')$.

Lemma 2.40. *The stratification $(G_{\mathrm{III}})_{T/} \to (\mathcal{S}_{\mathrm{III}})_{T/}$ is cell-like.*

Proof. Express the underlying tree T as the disjoint union of $T_i \in \mathcal{S}_{\mathrm{II}}$, so we have:

$$(G_{\mathrm{III}})_{T/} = \prod_{i \in I} (G_{\mathrm{II}})_{T_i/} \times \begin{cases} [0,1) & \mathfrak{s}(T) = \{0\} \\ (0,1) & \mathfrak{s}(T) = (0,1) \\ (0,1] & \mathfrak{s}(T) = \{1\} \end{cases} \tag{2.37}$$

$$(\mathcal{S}_{\mathrm{III}})_{T/} = \prod_{i \in I} (\mathcal{S}_{\mathrm{II}})_{T_i/} \times \begin{cases} \{\{0\} < (0,1)\} & \mathfrak{s}(T) = \{0\} \\ \{(0,1)\} & \mathfrak{s}(T) = (0,1) \\ \{(0,1) > \{1\}\} & \mathfrak{s}(T) = \{1\} \end{cases} \tag{2.38}$$

Now apply Lemma 2.38 and note that the product of two cell-like stratifications is again cell-like. ☐

Definition 2.41 (Space G_{IV})**.** Let $T \in \mathcal{S}_{\mathrm{IV}}$. For $\mathfrak{s}(T) \in \{\{0\}, (0,\infty)\}$, we define:

$$(G_{\mathrm{IV}})_{T/} := \Big\{(\{g_e\}_e, \{g_v\}_v) \in (0,\infty]^{E^{\mathrm{int}}(T)} \times (0,\infty]^{V_{00}(T)} \;\Big|\; g_v = g_e + g_{v'} \text{ for } v \xrightarrow{e} v' \text{ and } v \in V_{00}(T)\Big\} \times \begin{cases} [0,\infty) & \mathfrak{s}(T) = \{0\} \\ (0,\infty) & \mathfrak{s}(T) = (0,\infty) \end{cases} \tag{2.39}$$

(the first factor is identical to (2.31)); for $\mathfrak{s}(T) = \{\infty\}$ we define:

$$(G_{\mathrm{IV}})_{T/} := \left\{(\{g_e\}_e, \{g_v\}_v, t) \in (0,\infty]^{E^{\mathrm{int}}(T)} \times (0,\infty]^{V_{00}(T) \sqcup V_{11}(T)} \times (0,\infty] \;\middle|\; \begin{array}{r} g_v = g_e + g_{v'} \text{ for } v \xrightarrow{e} v',\ *(e) = 0,\ *(v) = 00 \\ g_v = g_e + g_{v'} \text{ for } v \xrightarrow{e} v',\ *(e) = 1,\ *(v) = 11 \\ t = g_e + g_{v'} \text{ for } v \xrightarrow{e} v',\ *(e) = 1,\ *(v) = 01 \end{array}\right\} \tag{2.40}$$

We interpret $g_{v'} = 0$ if it is undefined.

There is a natural stratification $(G_{\mathrm{IV}})_{T/} \to (\mathcal{S}_{\mathrm{IV}})_{T/}$, sending $g = (\{g_e\}, \{g_v\}, t)$ to the map $T \to T'$ which contracts those edges e for which $g_e < \infty$, increments $*^{-}(v)$ for those vertices v with $g_v < \infty$, and has $t \in \mathfrak{s}(T')$.

Lemma 2.42. *The stratification $(G_{\mathrm{IV}})_{T/} \to (\mathcal{S}_{\mathrm{IV}})_{T/}$ is cell-like.*

Proof. For $\mathfrak{s} \neq \{\infty\}$, this is just Lemma 2.40.

For $\mathfrak{s} = \{\infty\}$, argue as follows. For any given $f : T \to T'$, we may express $(G_{\mathrm{IV}})_{T//T'} \to (\mathcal{S}_{\mathrm{IV}})_{T//T'}$ as in (2.32)–(2.33), by expressing T' as a concatenation of maximal trees. Thus by induction, it suffices to show that $(G_{\mathrm{IV}})_{T/}$ is a topological manifold with boundary whose interior is the stratum corresponding to the maximal contraction of T. To see this, we do the same change of variables as in Lemma 2.38 to write:

$$(G_{\mathrm{IV}})_{T/} = \begin{Bmatrix} h_{v^{\mathrm{top}}} \in [0,\infty) & \text{if } *(v^{\mathrm{top}}) = 00 \\ q_e \in (-\infty,\infty) & \text{for } v \xrightarrow{e} v' \text{ with } *(v') = 00 \\ \tau \in [0,\infty) & \\ q_e \in (-\infty,\infty) & \text{for } v \xrightarrow{e} v' \text{ with } *(v') = 11 \\ h_e \in [0,\infty) & \text{for } *(e) = 2 \end{Bmatrix} \tag{2.41}$$

This is clearly a topological manifold of dimension $\operatorname{codim} T + 1$, and its maximal stratum is the locus where $\tau > 0$, $h_{v^{\mathrm{top}}} > 0$, and $h_e > 0$, which is clearly its interior. $\square$

3 Implicit atlases

In this section, we define (topological) implicit atlases with boundary and cell-like stratification (in the sense of [Par15, §§3,6]) on the moduli spaces $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$ stratified by $\mathcal{S}_{\mathrm{I}}$, $\mathcal{S}_{\mathrm{II}}$, $\mathcal{S}_{\mathrm{III}}$, $\mathcal{S}_{\mathrm{IV}}$. The construction of implicit atlases we give here follows the general procedure introduced in [Par15, §§1–2,9–10]. The basic point is to define appropriate thickened moduli spaces and to check that their regular loci cover the original moduli spaces.

3.1 Sets of thickening datums A_{I}, A_{II}, A_{III}, A_{IV}

We first define sets of thickening datums A_{I}, A_{II}, A_{III}, A_{IV} for the implicit atlases on $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$. Roughly speaking, a thickening datum α is a collection of data $(r_\alpha, D_\alpha, E_\alpha, \lambda_\alpha)$ from which we will construct a "thickened" version of a given moduli space by: (1) adding r_α marked points to the domain (constrainted to lie on the divisor $\hat{D}_\alpha$), (2) adding an extra parameter e_α lying in the vector space E_α, and (3) adding an extra term $\lambda_\alpha(e_\alpha)$ (which depends on the location of the r_α added points and the positive/negative ends $S = E^\pm(T)$ in the domain) to the pseudo-holomorphic curve equation.

Recall that $\overline{\mathcal{M}}_{0,n}$ ($n \geq 3$) denotes the Deligne–Mumford moduli space of stable nodal Riemann surfaces of genus zero with n marked points labeled with $\{1,\ldots,n\}$. We denote by $\overline{\mathcal{C}}_{0,n} \to \overline{\mathcal{M}}_{0,n}$ the universal family. Recall that $\overline{\mathcal{M}}_{0,n}$ is a compact smooth manifold. We usually prefer to label the marked points using a set other than $\{1,\ldots,n\}$, so we will also use the notation $\overline{\mathcal{M}}_{0,n}$ and $\overline{\mathcal{C}}_{0,n}$ when n is a *finite set* ($\#n \geq 3$) used to label the marked points.

Definition 3.1 (Set of thickening datums A_{I})**.** A *thickening datum* α for data as in Setup I along with a finite set S consists of the following data:

i. $r_\alpha \geq 0$ an integer such that $r_\alpha + \#S \geq 3$.

ii. E_α a finite-dimensional real vector space equipped with an action of S_{r_α} and an isomorphism $E_\alpha \xrightarrow{\sim} \mathbb{R}^{\dim E_\alpha}$.

iii. $D_\alpha \subseteq Y$ a compact codimension two submanifold with boundary. We let $\hat{D}_\alpha := \mathbb{R} \times D_\alpha \subseteq \hat{Y}$.

iv. $\lambda_\alpha : E_\alpha \to C^\infty(\hat{Y} \times \overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}, T\hat{Y} \otimes_{\mathbb{R}} \Omega^{0,1}_{\overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}/\overline{\mathcal{M}}_{0,S\cup\{1,\ldots,r_\alpha\}}})^{\mathbb{R}}$ an S_{r_α}-equivariant linear map vanishing in a neighborhood of the nodes and S-marked points of the fibers of $\overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}} \to \overline{\mathcal{M}}_{0,S\cup\{1,\ldots,r_\alpha\}}$. The superscript $^{\mathbb{R}}$ indicates taking the subspace of $\mathbb{R}$-invariant sections (where $\mathbb{R}$ acts on $\hat{Y}$ by translation).

We denote by $A_{\mathrm{I}}(S)$ the set[10] of such thickening datums.

Definition 3.2 (Set of thickening datums A_{II})**.** A *thickening datum* α for data as in Setup II along with a finite set S consists of the following data:

i. r_α, E_α as in Definition 3.1(i),(ii).

ii. $D_\alpha^\pm \subseteq Y^\pm$, $\lambda_\alpha^\pm : E_\alpha \to C^\infty(\hat{Y}^\pm \times \overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}, T\hat{Y}^\pm \otimes_{\mathbb{R}} \Omega^{0,1}_{\overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}/\overline{\mathcal{M}}_{0,S\cup\{1,\ldots,r_\alpha\}}})^{\mathbb{R}}$ as in Definition 3.1(iii),(iv).

iii. $\hat{D}_\alpha \subseteq \hat{X}$ a closed codimension two submanifold with boundary. We require that $\hat{D}_\alpha$ coincide (via (1.9)–(1.10)) with $\hat{D}_\alpha^\pm$ outside a compact subset of $\hat{X}$.

iv. $\lambda_\alpha : E_\alpha \to C^\infty(\hat{X} \times \overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}, T\hat{X} \otimes_{\mathbb{R}} \Omega^{0,1}_{\overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}/\overline{\mathcal{M}}_{0,S\cup\{1,\ldots,r_\alpha\}}})$ an S_{r_α}-equivariant linear map vanishing in a neighborhood of the nodes and S-marked points. We require that λ_α coincide (via (1.9)–(1.10)) with $\lambda_\alpha^\pm$ outside a compact subset of $\hat{X}$.

We denote by $A_{\mathrm{II}}(S)$ the set of such thickening datums.

Definition 3.3 (Set of thickening datums A_{III})**.** A *thickening datum* α for data as in Setup III is identical to a thickening datum for Setup II. This makes sense since in Setup III, the identifications (1.9)–(1.10) are independent of t; note also that the definition of a thickening datum does not make reference to λ^t or J^t. We denote by $A_{\mathrm{III}}(S)$ the set of such thickening datums.

Definition 3.4 (Set of thickening datums A_{IV})**.** A *thickening datum* α for data as in Setup IV along with a finite set S consists of the following data:

i. r_α, E_α as in Definition 3.1(i),(ii).

ii. $\{D_\alpha^i \subseteq Y^i\}_{i=0,1,2}$, $\{\lambda_\alpha^i : E_\alpha \to C^\infty(\hat{Y}^i \times \overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}, T\hat{Y}^i \otimes_{\mathbb{R}} \Omega^{0,1}_{\overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}/\overline{\mathcal{M}}_{0,S\cup\{1,\ldots,r_\alpha\}}})^{\mathbb{R}}\}_{i=0,1,2}$ as in Definition 3.1(iii),(iv).

iii. $\{\hat{D}_\alpha^{i,i+1} \subseteq \hat{X}^{i,i+1}\}_{i=0,1}$ as in Definition 3.2(iii).

iv. $\hat{D}_\alpha^{02,t} \subseteq \hat{X}^{02,t}$ a smoothly varying family of divisors for $t \in [0,\infty)$ as in Definition 3.2(iii) which coincides with the descent of $\hat{D}_\alpha^{01} \sqcup \hat{D}_\alpha^{12}$ for sufficiently large t. "Smoothly varying" means the projection to $[0,\infty)$ is a submersion (and remains a submersion when restricted to the boundary).

v. $\{\lambda_\alpha^{i,i+1} : E_\alpha \to C^\infty(\hat{X}^{i,i+1} \times \overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}, T\hat{X}^{i,i+1} \otimes_{\mathbb{R}} \Omega^{0,1}_{\overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}/\overline{\mathcal{M}}_{0,S\cup\{1,\ldots,r_\alpha\}}})\}_{i=0,1}$ as in Definition 3.2(iv).

vi. $\lambda_\alpha^{02,t} : E_\alpha \to C^\infty(\hat{X}^{02,t} \times \overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}, T\hat{X}^{02,t}, \otimes_{\mathbb{R}} \Omega^{0,1}_{\overline{\mathcal{C}}_{0,S\cup\{1,\ldots,r_\alpha\}}/\overline{\mathcal{M}}_{0,S\cup\{1,\ldots,r_\alpha\}}})$ as in Definition 3.2(iv) coinciding with the descent of $\lambda_\alpha^{01} \sqcup \lambda_\alpha^{12}$ for sufficiently large t, and varying smoothly with t.

[10]This is a set due to the choice of isomorphism $E_\alpha \xrightarrow{\sim} \mathbb{R}^{\dim E_\alpha}$; indeed this is the only purpose for including this choice of isomorphism in the definition.

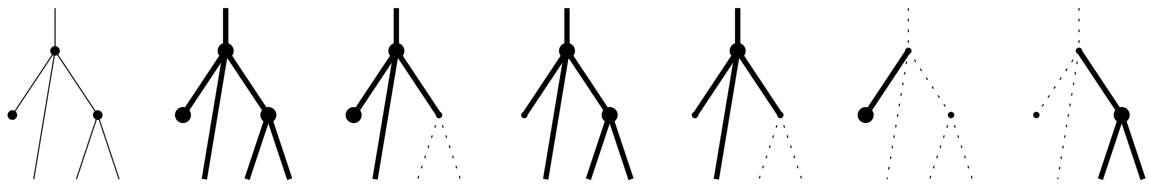

Figure 7: A tree and its six subtrees as in Definition 3.5.

We denote by $A_{\mathrm{IV}}(S)$ the set of such thickening datums.

3.2 Index sets $\bar{A}_{\mathrm{I}}$, $\bar{A}_{\mathrm{II}}$, $\bar{A}_{\mathrm{III}}$, $\bar{A}_{\mathrm{IV}}$

We now define the index sets $\bar{A}_{\mathrm{I}}$, $\bar{A}_{\mathrm{II}}$, $\bar{A}_{\mathrm{III}}$, $\bar{A}_{\mathrm{IV}}$ of the implicit atlases on $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$ as unions of copies of the sets of thickening datums A_{I}, A_{II}, A_{III}, A_{IV}.

We will use the short-hand $A_*(T)$ to mean $A_*(E^{\pm}(T))$ for $T \in \mathcal{S}_*$; note that a map $T \to T'$ induces an identification $A_*(T) = A_*(T')$.

Definition 3.5 (Subtree)**.** A *subtree* $T' \subseteq T$ shall mean one obtained by choosing some subset of the vertices to keep $V(T') \subseteq V(T)$, and keeping all the edges which are adjacent to at least one of these vertices; a subtree is also required to be connected and non-empty (see Figure 7).

Definition 3.6 (Index set $\bar{A}_{\mathrm{I}}$)**.** For $T \in \mathcal{S}_{\mathrm{I}}$, we define:

$$\bar{A}_{\mathrm{I}}(T) := \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{I}}} A_{\mathrm{I}}(T') \tag{3.1}$$

(union over subtrees). A subtree $T \supseteq T' \in \mathcal{S}_{\mathrm{I}}$ is means a subtree as in Definition 3.5, with decorations inherited from those of T (it will not matter how we choose basepoints for subtrees).

Definition 3.7 (Index set $\bar{A}_{\mathrm{II}}$)**.** For $T \in \mathcal{S}_{\mathrm{II}}$, we define:

$$\bar{A}_{\mathrm{II}}(T) := \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^+} A_{\mathrm{I}}^+(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^-} A_{\mathrm{I}}^-(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{II}}} A_{\mathrm{II}}(T') \tag{3.2}$$

More precisely, the unions are over the following types of subtrees:
i. $T \supseteq T' \in \mathcal{S}_{\mathrm{II}}$, i.e. those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^{\pm} \in E^{\pm}(T')$.
ii. $T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^+$, i.e. those for which all edges and vertices have $* = 0$.
iii. $T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^-$, i.e. those for which all edges and vertices have $* = 1$.
(where $\mathcal{S}_{\mathrm{I}}^{\pm} = \mathcal{S}_{\mathrm{I}}(Y^{\pm}, \lambda^{\pm})$). Note that a given physical subtree $T' \subseteq T$ may appear in more than one $\bigsqcup$ above.

Definition 3.8 (Index set $\bar{A}_{\mathrm{III}}$)**.** For $T \in \mathcal{S}_{\mathrm{III}}$, we define:

$$\begin{aligned}\bar{A}_{\mathrm{III}}(T) := \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^+} A_{\mathrm{I}}^+(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^-} A_{\mathrm{I}}^-(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{t=0}} A_{\mathrm{II}}^{t=0}(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{t=1}} A_{\mathrm{II}}^{t=1}(T') \\ \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{III}}} A_{\mathrm{III}}(T')\end{aligned} \tag{3.3}$$

More precisely, the unions are over the following types of subtrees:

i. $T \supseteq T' \in \mathcal{S}_{\mathrm{III}}$, i.e. those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^\pm \in E^\pm(T')$.
ii. $T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{t=0}$, i.e. those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^\pm \in E^\pm(T')$ if $\mathfrak{s}(T) = \{0\}$.
iii. $T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{t=1}$, i.e. those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^\pm \in E^\pm(T')$ if $\mathfrak{s}(T) = \{1\}$.
iv. $T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^{+}$, i.e. those for which all edges and vertices have $* = 0$.
v. $T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^{-}$, i.e. those for which all edges and vertices have $* = 1$.

Definition 3.9 (Index set $\bar{A}_{\mathrm{IV}}$)**.** For $T \in \mathcal{S}_{\mathrm{IV}}$, we define:

$$\begin{aligned}\bar{A}_{\mathrm{IV}}(T) := \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^{0}} A_{\mathrm{I}}^{0}(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^{2}} A_{\mathrm{I}}^{2}(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{02}} A_{\mathrm{II}}^{02}(T') \\ \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{01}} A_{\mathrm{II}}^{01}(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^{1}} A_{\mathrm{I}}^{1}(T') \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{12}} A_{\mathrm{II}}^{12}(T') \\ \sqcup \bigsqcup_{T \supseteq T' \in \mathcal{S}_{\mathrm{IV}}} A_{\mathrm{IV}}(T')\end{aligned} \tag{3.4}$$

More precisely, the unions are over the following types of subtrees:
i. $T \supseteq T' \in \mathcal{S}_{\mathrm{IV}}$, i.e. those with $*(e^+) = 0$ and $*(e^-) = 2$ for $e^\pm \in E^\pm(T')$.
ii. $T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{01}$, i.e. those with $*(e^+) = 0$ and $*(e^-) = 1$ for $e^\pm \in E^\pm(T')$ and $\mathfrak{s}(T) = \{\infty\}$.
iii. $T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{12}$, i.e. those with $*(e^+) = 1$ and $*(e^-) = 2$ for $e^\pm \in E^\pm(T')$ and $\mathfrak{s}(T) = \{\infty\}$.
iv. $T \supseteq T' \in \mathcal{S}_{\mathrm{II}}^{02}$, i.e. those with $*(e^+) = 0$ and $*(e^-) = 2$ for $e^\pm \in E^\pm(T')$ and $\mathfrak{s}(T) = \{0\}$.
v. $T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^{0}$, i.e. those for which all edges and vertices have $* = 0$.
vi. $T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^{1}$, i.e. those for which all edges and vertices have $* = 1$.
vii. $T \supseteq T' \in \mathcal{S}_{\mathrm{I}}^{2}$, i.e. those for which all edges and vertices have $* = 2$.

A morphism $T \to T'$ induces a natural inclusion:

$$\bar{A}_*(T') \hookrightarrow \bar{A}_*(T) \tag{3.5}$$

(since given $T \to T'$, any subtree $T'' \subseteq T'$ pulls back to a subtree of T).

For any concatenation $\{T_i\}_i$, there is a natural inclusion:

$$\bigsqcup_i \bar{A}_*(T_i) \hookrightarrow \bar{A}_*(\#_i T_i) \tag{3.6}$$

3.3 Thickened moduli spaces

We now define the thickened moduli spaces for the implicit atlases on $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$.

Definition 3.10 (Moduli space $\mathcal{M}_{\mathrm{I}}(T)_I$)**.** Let $T \in \mathcal{S}_{\mathrm{I}}$ and let $I \subseteq \bar{A}_{\mathrm{I}}(T)$ be finite. An *I-thickened pseudo-holomorphic building of type T* consists of the following data:
i. C_v and $p_{v,e}$ as in Definition 2.9(i). For $\alpha \in I$, let $C_\alpha := \bigsqcup_{v \in T_\alpha} C_v/\sim$, where $T \supseteq T_\alpha \in \mathcal{S}_{\mathrm{I}}$ denotes the subtree indexing the term in (3.1) containing $\alpha \in I$, and $\sim$ identifies $p_{v,e} \sim p_{v',e}$ for interior edges $v \xrightarrow{e} v'$ of T_α (see Figure 8).
ii. u_v, b_e, m_e as in Definition 2.9(ii),(iii),(iv),(v).
iii. For all $\alpha \in I$, we require that $(u|C_\alpha) \pitchfork \hat{D}_\alpha$ with exactly r_α intersections, which together with $\{p_{v,e}\}_{v,e}$ stabilize C_α. By $(u|C_\alpha) \pitchfork \hat{D}_\alpha$, we mean that $\gamma_e \cap D_\alpha = \varnothing$ for edges $e \in E(T_\alpha)$, $(u|C_\alpha)^{-1}(\partial\hat{D}_\alpha) = \varnothing$, $(u|C_\alpha)^{-1}(\hat{D}_\alpha)$ does not contain any node, and $(du)_p : T_pC_\alpha \to T_{u(p)}\hat{Y}/T_{u(p)}\hat{D}_\alpha$ is surjective for $p \in (u|C_\alpha)^{-1}(\hat{D}_\alpha)$.

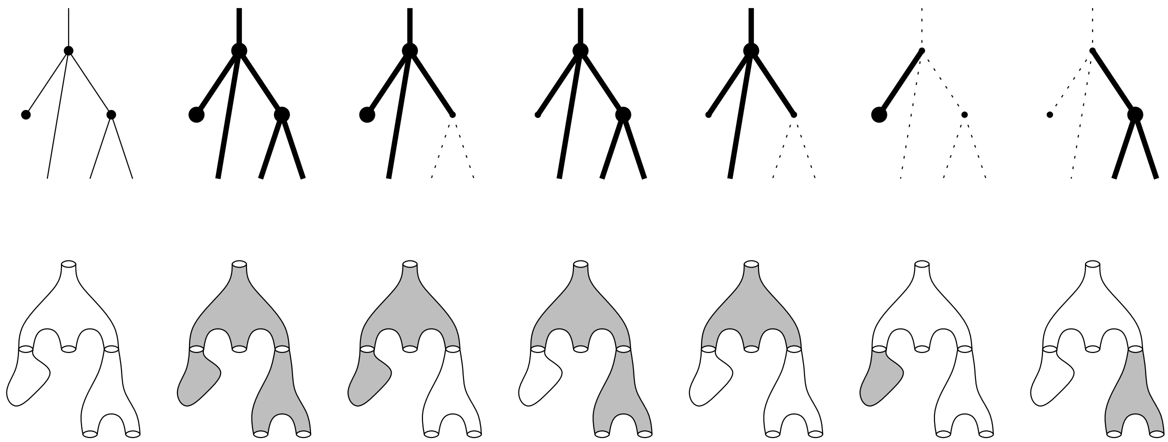

Figure 8: A tree with its six subtrees, and a corresponding pseudo-holomorphic building with its corresponding subbuildings.

iv. $\{\phi_\alpha : C_\alpha \to \overline{\mathcal{C}}_{0,E^\pm(T_\alpha)\cup\{1,\ldots,r_\alpha\}}\}_{\alpha\in I}$, where each ϕ_α maps C_α isomorphically onto a fiber of $\overline{\mathcal{C}}_{0,E^\pm(T_\alpha)\cup\{1,\ldots,r_\alpha\}}$, where C_α is equipped with its given marked points $p_{v,e}$ for $e \in E^\pm(T_\alpha)$ and any marking of $(u|C_\alpha)^{-1}(\hat{D}_\alpha)$ with $\{1,\ldots,r_\alpha\}$. Note that under (iii) above, choosing ϕ_α is equivalent to choosing a marking of $(u|C_\alpha)^{-1}(\hat{D}_\alpha)$ with $\{1,\ldots,r_\alpha\}$.

v. $\{e_\alpha \in E_\alpha\}_{\alpha\in I}$.

vi. We require that u satisfy the following I-thickened pseudo-holomorphic curve equation:

$$\Big(du + \sum_{\alpha\in I} \lambda_\alpha(e_\alpha)(\phi_\alpha(\cdot), u(\cdot))\Big)^{0,1}_{\hat{J}} = 0 \tag{3.7}$$

Note that the term in $\sum_{\alpha\in I}$ corresponding to α makes sense only over C_α, and we define it to be zero elsewhere.

An *isomorphism* between I-thickened pseudo-holomorphic buildings of type T is defined as in Definition 2.9, with the additional requirements that $e_\alpha = e'_\alpha$ and $\phi_{\alpha,v} = \phi'_{\alpha,v} \circ i_v$ for $v \in T_\alpha$. We denote by $\mathcal{M}_{\mathrm{I}}(T)_I$ the set of isomorphism classes of stable I-thickened pseudo-holomorphic buildings of type T.

Note that the sum over α in (3.7) is supported away from the punctures $p_{v,e} \in C_v$, and hence u_v is genuinely $\hat{J}$-holomorphic near $p_{v,e}$. Note also that (3.7) is equivalent to the assertion that the graph $(\mathrm{id}, u_v) : C_v \to C_v \times \hat{Y}$ is pseudo-holomorphic for the almost complex structure on $C_v \times \hat{Y}$ given by:

$$\begin{pmatrix} j_{C_v} & 0 \\ (\sum_{\alpha\in I} \lambda(e_\alpha)(\phi_\alpha(\cdot), \cdot))^{0,1} & \hat{J} \end{pmatrix} \tag{3.8}$$

Hence solutions to the I-thickened pseudo-holomorphic curve equation enjoy all of the nice elliptic estimates which apply to solutions to the (usual) pseudo-holomorphic curve equation.

Definition 3.11 (Moduli space $\mathcal{M}_{\mathrm{II}}(T)_I$)**.** Let $T \in \mathcal{S}_{\mathrm{II}}$ and let $I \subseteq \bar{A}_{\mathrm{II}}(T)$. An *$I$-thickened pseudo-holomorphic building of type* T consists of the same data as in Definition 3.10 (the

only difference being the target of u_v and using $\hat{D}^\pm_\alpha, D_\alpha$ and $\lambda^\pm_\alpha, \lambda_\alpha$ appropriately). We denote by $\mathcal{M}_{\mathrm{II}}(T)_I$ the set of isomorphism classes of stable I-thickened pseudo-holomorphic buildings of type T.

Definition 3.12 (Moduli space $\mathcal{M}_{\mathrm{III}}(T)_I$)**.** Let $T \in \mathcal{S}_{\mathrm{III}}$ and let $I \subseteq \bar{A}_{\mathrm{III}}(T)$. We denote by $\mathcal{M}_{\mathrm{III}}(T)_I$ the union over $t \in \mathfrak{s}(T)$ of the set of isomorphism classes of stable I-thickened pseudo-holomorphic buildings of type T.

Definition 3.13 (Moduli space $\mathcal{M}_{\mathrm{IV}}(T)_I$)**.** Let $T \in \mathcal{S}_{\mathrm{IV}}$ and let $I \subseteq \bar{A}_{\mathrm{IV}}(T)$. We denote by $\mathcal{M}_{\mathrm{IV}}(T)_I$ the union over $t \in \mathfrak{s}(T)$ of the set of isomorphism classes of stable I-thickened pseudo-holomorphic buildings of type T.

Definition 3.14 (Moduli spaces $(\overline{\mathcal{M}}_{\mathrm{I}})_I$, $(\overline{\mathcal{M}}_{\mathrm{II}})_I$, $(\overline{\mathcal{M}}_{\mathrm{III}})_I$, $(\overline{\mathcal{M}}_{\mathrm{IV}})_I$)**.** For $T \in \mathcal{S}_*$ and $I \subseteq \bar{A}_*(T)$, we define:

$$\overline{\mathcal{M}}_*(T)_I := \bigsqcup_{\substack{T' \to T \\ \overline{\mathcal{M}}_*(T') \neq \varnothing}} \mathcal{M}_*(T')_I / \operatorname{Aut}(T'/T) \tag{3.9}$$

Each such set $(\overline{\mathcal{M}}_*)_I$ has a natural Gromov topology which is Hausdorff.

The stratifications (2.5) are clearly defined on the thickened moduli spaces $(\overline{\mathcal{M}}_{\mathrm{I}})_I$, $(\overline{\mathcal{M}}_{\mathrm{II}})_I$, $(\overline{\mathcal{M}}_{\mathrm{III}})_I$, $(\overline{\mathcal{M}}_{\mathrm{IV}})_I$. The tautological functorial structure (2.3) (combined with (3.5)) and (2.4) (combined with (3.6)) also exists for the thickened moduli spaces.

There are natural maps $s_\alpha : (\overline{\mathcal{M}}_*)_I \to E_\alpha$ for $\alpha \in I$, which are $\Gamma_\alpha := S_{r_\alpha}$-equivariant for the natural action of Γ_α on $(\overline{\mathcal{M}}_*)_I$ (by acting on e_α and ϕ_α). There are natural forgetful maps $\psi_{IJ} : (s_{J\setminus I}|(\overline{\mathcal{M}}_*)_J)^{-1}(0)/\Gamma_{J\setminus I} \to (\overline{\mathcal{M}}_*)_I$. Each ψ_{IJ} is a bijection whose image $U_{IJ} \subseteq (\overline{\mathcal{M}}_*)_I$ is open (as it is the locus of points satisfying the transversality condition in Definition 3.10(iii) for $\alpha \in J \setminus I$). Inspection of the definition of the Gromov topology shows that ψ_{IJ} is in fact a homeomorphism. Clearly $(\overline{\mathcal{M}}_*)_\varnothing = \overline{\mathcal{M}}_*$.

3.4 Linearized operators

We now describe the linearized operators associated to I-thickened pseudo-holomorphic buildings.

Definition 3.15 (Linearized operator)**.** Given an I-thickened pseudo-holomorphic building of type T, there is an associated linearized operator:

$$E_I \oplus \bigoplus_{v \in V(T)} \tilde{W}^{k,2,\delta}(C_v, u_v^* T\hat{X}_v) \to \bigoplus_{v \in V(T)} W^{k-1,2,\delta}(\tilde{C}_v, u_v^*(T\hat{X}_v)_{\hat{J}_v} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}_v}) \tag{3.10}$$

A point in a moduli space $(\overline{\mathcal{M}}_*)_I$ is called *regular* iff the linearized operator (3.10) of the corresponding pseudo-holomorphic building is surjective (it follows from elliptic regularity theory that this condition is independent of k and δ). Let $(\overline{\mathcal{M}}_*)_I^{\mathrm{reg}} \subseteq (\overline{\mathcal{M}}_*)_I$ denote the locus of points which are regular.

3.5 Stabilization of pseudo-holomorphic curves with divisors

We now verify the covering axiom for the implicit atlases which we are in the process of defining. Namely, we show that the moduli spaces $\overline{\mathcal{M}}_*$ are covered by the regular loci in their thickenings $(\overline{\mathcal{M}}_*)_I^{\mathrm{reg}}$. The essential content is to show that at every point in each moduli space $\overline{\mathcal{M}}_*$, there exists a divisor (in the sense of Definition 3.1(iii)) which stabilizes the domain (i.e. satisfies Definition 3.10(iii)).

Lemma 3.16. *Let $u : D^2 \to (X, J)$ be J-holomorphic (for an almost complex manifold (X, J)). Then either $du : T_pC \to T_{u(p)}X$ is injective for some $p \in D^2$ or*[11] *u is constant.*

Proof. If du is non-injective, it must be zero by J-holomorphicity. ☐

Lemma 3.17. *Let $u : D^2 \to (\hat{Y}, \hat{J})$ be $\hat{J}$-holomorphic (for data as in Setup I). Denote by $\pi_\xi : T\hat{Y} \to \xi$ the projection under the splitting $T\hat{Y} = \xi \oplus \mathbb{R}R_\lambda \oplus \mathbb{R}\partial_s$. Then either $\pi_\xi du : T_pC \to \xi_{u(p)}$ is injective for some $p \in D^2$ or*[12] *u factors through $\mathrm{id} \times \gamma : \mathbb{R} \times \mathbb{R} \to \mathbb{R} \times Y$ for some Reeb trajectory $\gamma : \mathbb{R} \to Y$.*

Proof. If $\pi_\xi du$ is non-injective, it must be zero. If $\pi_\xi du$ vanishes identically, then du is everywhere tangent to the 2-dimensional foliation of $\hat{Y}$ by $\mathbb{R}R_\lambda \oplus \mathbb{R}\partial_s$, and thus u factors through one of its leaves. ☐

Proposition 3.18. *For every $x \in \overline{\mathcal{M}}_{\mathrm{I}}(T)$, there exists $\alpha \in A_{\mathrm{I}}(T)$ such that $x \in U_{\varnothing,\{\alpha\}}$ and $\psi_{\varnothing,\{\alpha\}}^{-1}(x) \subseteq \overline{\mathcal{M}}_{\mathrm{I}}(T)_{\{\alpha\}}^{\mathrm{reg}}$.*

Proof. The point x is an isomorphism class of stable pseudo-holomorphic building of type $T' \to T$.

We claim that for all $v \in V(T')$, either C_v is stable (i.e. the degree of v is ≥ 3) or $\pi_\xi du_v$ is injective somewhere on C_v. To see this, suppose that $\pi_\xi du_v \equiv 0$ and apply Lemma 3.17 to $u_v : C_v \to \hat{Y}$. If the resulting Reeb trajectory $\gamma : \mathbb{R} \to Y$ is not a closed orbit, then u_v factors through $\mathbb{R} \times \mathbb{R} \to \mathbb{R} \times Y$, and consideration of the positive puncture of C_v leads to a contradiction. Thus u_v factors through a trivial cylinder $\mathbb{R} \times S^1 \to \mathbb{R} \times Y$ for some simple Reeb orbit $\gamma : S^1 \to Y$. The map $C_v \to \mathbb{R} \times S^1$ is holomorphic, and it must have ramification points, as otherwise the building x would be unstable. It now follows from Riemann–Hurwitz that C_v is stable. Thus the claim is valid.

Now using the claim, it follows from Sard's theorem that there exists $D_\alpha \subseteq Y$ satisfying Definition 3.10(iii) for some $r_\alpha \geq 0$. Now to show the existence of E_α and λ_α so that x is regular for the thickening datum $\alpha = (D_\alpha, r_\alpha, E_\alpha, \lambda_\alpha)$, it suffices to show that:

$$C_c^\infty(\tilde{C}_v \setminus (\{p_{v,e}\}_e \cup N_v), u_v^* T\hat{Y}_{\hat{J}} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}_v}) \tag{3.11}$$

surjects onto the (finite-dimensional) cokernel of the linearized operator at x.

Now (3.11) is dense in $W^{k-1,2,\delta}$ for $k = 2$ and admissible $\delta > 0$, so we have the desired surjectivity. One could also use Sobolev spaces with weights near the nodes (see Definition 5.2.1) and then (3.11) is dense for all $k \geq 2$ and admissible $\delta > 0$. ☐

[11]In fact, it is a standard (but nontrivial) fact that the first alternative can be strengthened to state that the zeroes of du form a discrete set (see [MS04, Lemma 2.4.1]).

[12]In fact, it is a standard (but nontrivial) fact that the first alternative can be strengthened to state that the zeroes of $\pi_\xi du$ form a discrete set (see Hofer–Wysocki–Zehnder [HWZ95, Proposition 4.1]).

Remark 3.19. Let us give an alternative argument that (3.11) surjects onto the cokernel of the linearized operator D. It suffices to show that there is no continuous linear functional on the cokernel of D which vanishes on (3.11). Such a continuous linear functional would be a distribution ϵ valued in $u_v^*T\hat{Y}_{\hat{\jmath}} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}_v}$ supported over the pre-images of the nodes $\tilde{N}_v \subseteq \tilde{C}_v$, with the property that $\langle \epsilon, D\xi \rangle = 0$ for all smooth (compactly supported) $\xi : C_v \to u_v^*T\hat{Y}$. Now this implies that D^* (the formal adjoint) applied to ϵ is a linear combination of δ-functions supported over $\tilde{N}_v$. But such δ-functions live in H^{-2} and the formal adjoint is elliptic, so this means $\epsilon \in H^{-1}$, which contains no nonzero distributions supported at single points.

Proposition 3.20. *For every $x \in \overline{\mathcal{M}}_{\mathrm{II}}(T)$, there exists $\alpha \in A_{\mathrm{II}}(T)$ such that $x \in U_{\varnothing,\{\alpha\}}$ and $\psi^{-1}_{\varnothing,\{\alpha\}}(x) \subseteq \overline{\mathcal{M}}_{\mathrm{II}}(T)^{\mathrm{reg}}_{\{\alpha\}}$.*

Proof. The point x is an isomorphism class of stable pseudo-holomorphic building of type $T' \to T$.

As in the proof of Proposition 3.18, for $v \in V(T')$ with $*(v) = 00$ or $*(v) = 11$, either C_v is stable or $\pi_\xi du_v$ is injective somewhere on C_v. For $*(v) = 01$, every irreducible component of C_v is either stable or has a point where du_v injective by Lemma 3.16.

Now we can find a (compact) divisor $\hat{D}_\alpha \subseteq \hat{X}$ such that $u \pitchfork \hat{D}_\alpha$ with intersections stabilizing every C_v with $*(v) = 01$. Now we consider the remaining unstable C_v ($*(v) = 00$ or $*(v) = 11$), and we choose divisors $D^\pm_\alpha \subseteq Y^\pm$ stabilizing these. We then cutoff $\hat{D}^\pm_\alpha$ near infinity in $\hat{X}$ and add this to $\hat{D}_\alpha$. Thus $u \pitchfork \hat{D}_\alpha$ with r_α intersections which stabilize C.

Now $(E_\alpha, \lambda_\alpha)$ are constructed as in Proposition 3.18. ☐

Proposition 3.21. *For every $x \in \overline{\mathcal{M}}_{\mathrm{III}}(T)$, there exist $\alpha_i \in A_{\mathrm{III}}(T_i)$ such that $x \in U_{\varnothing,\{\alpha_i\}_i}$ and $\psi^{-1}_{\varnothing,\{\alpha_i\}_i}(x) \subseteq \overline{\mathcal{M}}_{\mathrm{III}}(T)^{\mathrm{reg}}_{\{\alpha_i\}_i}$ (writing $T = \bigsqcup_i T_i$ with T_i connected and non-empty).*

Proof. Apply Proposition 3.20 to each subbuilding of type T_i to get α_i. ☐

Proposition 3.22. *For every $x \in \overline{\mathcal{M}}_{\mathrm{IV}}(T)$, there exist $\alpha_i \in A_{\mathrm{IV}}(T_i)$ such that $x \in U_{\varnothing,\{\alpha_i\}_i}$ and $\psi^{-1}_{\varnothing,\{\alpha_i\}_i}(x) \subseteq \overline{\mathcal{M}}_{\mathrm{IV}}(T)^{\mathrm{reg}}_{\{\alpha_i\}_i}$ (writing $T = \bigsqcup_i T_i$ with T_i connected and non-empty).*

Proof. Essentially the same as the proof of Propositions 3.20–3.21. ☐

3.6 Local structure of thickened moduli spaces via models G_{I}, G_{II}, G_{III}, G_{IV}

We now state the precise sense in which the spaces G_{I}, G_{II}, G_{III}, G_{IV} are local topological models for the regular loci in the thickened moduli spaces $(\overline{\mathcal{M}}_{\mathrm{I}})_I$, $(\overline{\mathcal{M}}_{\mathrm{II}})_I$, $(\overline{\mathcal{M}}_{\mathrm{III}})_I$, $(\overline{\mathcal{M}}_{\mathrm{IV}})_I$. This statement is in essence a gluing theorem, and its proof is given in §5. It implies that the submersion and openness axioms hold for the implicit atlases we have defined, and it also gives canonical isomorphisms $\mathfrak{o}_{\overline{\mathcal{M}}_*(T)} = \mathfrak{o}_T$.

The (Banach space) implicit function theorem implies that $\mathcal{M}_*(T)^{\mathrm{reg}}_I$ is a (smooth) manifold of dimension $\mu(T) - \#V_s(T) + \dim \mathfrak{s}(T) + \dim E_I$ over the locus without nodes.

Theorem 3.23 (Local structure of $\overline{\mathcal{M}}_*(T)_I^{\mathrm{reg}}$)**.** *Fix $* \in \{\mathrm{I}, \mathrm{II}, \mathrm{III}, \mathrm{IV}\}$. Let $I \subseteq J \subseteq \bar{A}_*(T)$. Let $x_0 \in \mathcal{M}_*(T')_J/\mathrm{Aut}(T'/T) \subseteq \overline{\mathcal{M}}_*(T)_J$ (some $T' \to T$) be such that $s_{J\setminus I}(x_0) = 0$ and $\psi_{IJ}(x_0) \in \overline{\mathcal{M}}_*(T)_I^{\mathrm{reg}}$. Then $\mu(T) - \#V_s(T') - 2\#N(x_0) + \dim E_I \geq 0$ and there is a germ of homeomorphism:*

$$\big((G_*)_{T'//T} \times E_{J\setminus I} \times \mathbb{C}^{N(x_0)} \times \mathbb{R}^{\mu(T)-\#V_s(T')-2\#N(x_0)+\dim E_I}, (0,0,0)\big) \to \big(\overline{\mathcal{M}}_*(T)_J, x_0\big) \quad (3.12)$$

whose image lands in $\overline{\mathcal{M}}_*(T)_J^{\mathrm{reg}}$ and which commutes with the maps from both sides to $(\mathcal{S}_*)_{T'//T} \times \mathfrak{s}(T) \times E_{J\setminus I}$ (as well as the stratifications by number of nodes).

Proof. See §§5.1–5.3. ☐

The (Banach space) implicit function theorem moreover gives a canonical identification:

$$\mathfrak{o}_{\overline{\mathcal{M}}_*(T)_I^{\mathrm{reg}}} = \mathfrak{o}_T \otimes \mathfrak{o}_{E_I} \quad (3.13)$$

More precisely, this identification is made over the locus in $\mathcal{M}_*(T)_I^{\mathrm{reg}}$ without nodes, and has a unique continuous extension to all of $\overline{\mathcal{M}}_*(T)_I^{\mathrm{reg}}$ by virtue of the local topological description in Theorem 3.23. This identification is easily seen to be compatible with ψ_{IJ} and with concatenations. It is also compatible with morphisms $T' \to T$ in the following precise sense:

Theorem 3.24 (Compatibility of the "analytic" and "geometric" maps on orientations)**.** *The following diagram commutes:*

$$\begin{tikzcd}
\mathfrak{o}_{\overline{\mathcal{M}}_*(T')_I^{\mathrm{reg}}} \otimes \mathfrak{o}_{\mathbb{R}}^{\otimes V_s(T')} \otimes \mathfrak{o}_{\mathfrak{s}(T')}^\vee \arrow[r, equal, "(3.13)"] \arrow[d] & \mathfrak{o}_{T'}^0 \otimes \mathfrak{o}_{E_I} \arrow[d] \\
\mathfrak{o}_{\overline{\mathcal{M}}_*(T)_I^{\mathrm{reg}}} \otimes \mathfrak{o}_{\mathbb{R}}^{\otimes V_s(T)} \otimes \mathfrak{o}_{\mathfrak{s}(T)}^\vee \arrow[r, equal, "(3.13)"] & \mathfrak{o}_{T}^0 \otimes \mathfrak{o}_{E_I}
\end{tikzcd} \quad (3.14)$$

where the maps are as follows. The left vertical map is the "geometric" map induced by the local topological structure of $\overline{\mathcal{M}}_*(T)_I^{\mathrm{reg}}$ coming from (3.12). The right vertical map is the "analytic" map defined earlier via the "kernel gluing" operation.

Proof. See §5.4. ☐

4 Virtual fundamental cycles

In this section, we prove Theorems I, II, III, IV as stated in the introduction. Specifically, these main results follow from combining Definitions 4.34, 4.35 and Lemmas 4.36, 4.37.

Our work in this section relies heavily on the framework introduced in [Par15], and we begin with a review of the machinery which we will need. We use $\mathbb{Q}$ coeffients throughout.

4.1 Review of the VFC package

In this subsection, we review the framework introduced in [Par15] for defining the virtual fundamental cycle of a space equipped with an implicit atlas. For motivation, we refer the reader to [Par15, §§1–2], and for complete definitions, we refer the reader to [Par15, §§4–6,A]; our arguments in the rest of this section are similar to, though logically independent of, those in [Par15, §7].

Let X be a (compact Hausdorff) space equipped with a finite implicit atlas with boundary A. The implicit atlas A induces an orientation sheaf $\mathfrak{o}_X$ over X (and we set $\mathfrak{o}_{X\,\mathrm{rel}\,\partial} := j_! j^* \mathfrak{o}_X$ for $j : X \setminus \partial X \hookrightarrow X$). The atlas A also induces "virtual cochain complexes" $C^\bullet_{\mathrm{vir}}(X; A)$ and $C^\bullet_{\mathrm{vir}}(X \operatorname{rel} \partial; A)$, along with a natural map:

$$C^{\bullet-1}_{\mathrm{vir}}(\partial X; A) \to C^\bullet_{\mathrm{vir}}(X \operatorname{rel} \partial; A) \tag{4.1}$$

whose mapping cone is denoted $C^\bullet_{\mathrm{vir}}(X; A) := [C^\bullet_{\mathrm{vir}}(\partial X; A) \to C^\bullet_{\mathrm{vir}}(X \operatorname{rel} \partial; A)]$. There are natural isomorphisms:

$$H^\bullet_{\mathrm{vir}}(X; A) = \check{H}^\bullet(X, \mathfrak{o}_X) \tag{4.2}$$
$$H^\bullet_{\mathrm{vir}}(X \operatorname{rel} \partial; A) = \check{H}^\bullet(X, \mathfrak{o}_{X\,\mathrm{rel}\,\partial}) \tag{4.3}$$

We denote by $C_\bullet(E; A) := C_{\dim E_A + \bullet}(E_A, E_A \setminus 0; \mathfrak{o}^\vee_{E_A})^{\Gamma_A}$; note that there is a canonical isomorphism $H_\bullet(E; A) = \mathbb{Q}$ (concentrated in degree zero). The virtual fundamental cycle of X is represented (on the chain level!) by a canonical pushforward map:

$$C^{d+\bullet}_{\mathrm{vir}}(X \operatorname{rel} \partial; A) \to C_{-\bullet}(E; A) \tag{4.4}$$

where d is the virtual dimension of A. For an inclusion of implicit atlases $A \subseteq A'$ on the same space X, there are canonical quasi-isomorphisms:

$$C^\bullet_{\mathrm{vir}}(X \operatorname{rel} \partial; A) \xrightarrow{\sim} C^\bullet_{\mathrm{vir}}(X \operatorname{rel} \partial; A') \tag{4.5}$$
$$C_\bullet(E; A) \xrightarrow{\sim} C_\bullet(E; A') \tag{4.6}$$

which compose as expected and are compatible with (4.4) (both on the chain level).[13]

For spaces X and Y with implicit atlases A and B, there are product maps:

$$C^\bullet_{\mathrm{vir}}(X \operatorname{rel} \partial; A) \otimes C^\bullet_{\mathrm{vir}}(Y \operatorname{rel} \partial; B) \to C^\bullet_{\mathrm{vir}}(X \times Y \operatorname{rel} \partial; A \sqcup B) \tag{4.7}$$

and these are compatible with (4.4) and (4.5)–(4.6).

Remark 4.1. Though not necessary for the present construction, let us remark that given (4.5)–(4.6) as defined above, we can actually define $C^\bullet_{\mathrm{vir}}(-; A)$ and $C^\bullet_{\mathrm{vir}}(E; A)$ for arbitrary A by taking the direct limit over finite subsets.

[13] The maps (4.5)–(4.6) are slight modifications of the originals [Par15, §4.2, eq. (4.2.14)–(4.2.15)], which carry an extra factor of $\otimes C_\bullet(E; A' \setminus A)$ on the left. Let $[E_\alpha] \in C_0(E; \alpha)$ be the fundamental cycle obtained by pulling back some $[\mathbb{R}^n] \in C_n(\mathbb{R}^n, \mathbb{R}^n \setminus 0; \mathfrak{o}^\vee_{\mathbb{R}^n})$ (fixed once and for all) under the specified isomorphism $E_\alpha \xrightarrow{\sim} \mathbb{R}^{\dim E_\alpha}$ and averaging over Γ_α (alternatively, we could modify the definition of a thickening datum to include the data of a fundamental cycle $[E_\alpha] \in C_0(E; \alpha)$). Now (4.5)–(4.6) are defined by pre-composing the maps in [Par15] with $\otimes \bigotimes_{\alpha \in A' \setminus A} [E_\alpha]$.

Remark 4.2. Our convention to keep track of signs is the following. Everything is $\mathbb{Z}/2$-graded, and $\otimes$ is always the super tensor product (namely, where the isomorphism $A\otimes B \xrightarrow{\sim} B\otimes A$ is given by $a\otimes b\mapsto(-1)^{|a||b|}b\otimes a$, where $(f\otimes g)(a\otimes b):=(-1)^{|g||a|}f(a)\otimes g(b)$). Complexes are $(\mathbb{Z},\mathbb{Z}/2)$-bigraded; differentials are always odd, chain maps are always even, and chain homotopies are always odd. For specificity, let us declare that $\operatorname{Hom}(A,B)\otimes A\to B$ be given by $f\otimes a\mapsto f(a)$ (though it won't really matter for our arguments).

We conclude by stating precisely some results which are implicit in [Par15, §6].

Lemma 4.3 (c.f. [Par15, §A])**.** *Let $\mathcal{F}^\bullet$, $\mathcal{G}^\bullet$, $\mathcal{H}^\bullet$ be pure homotopy $\mathcal{K}$-sheaves on spaces X, Y, $X\times Y$, respectively. Let $\mathcal{F}^\bullet(K)\otimes\mathcal{G}^\bullet(K')\to\mathcal{H}^\bullet(K\times K')$ be a collection of maps compatible with restriction. Then the following diagram commutes:*

$$\begin{tikzcd}
H^\bullet\mathcal{F}^\bullet(X)\otimes H^\bullet\mathcal{G}^\bullet(Y) \arrow[r] \arrow[d] & H^\bullet\mathcal{H}^\bullet(X\times Y) \arrow[d] \\
\check{H}^\bullet(X;H^0\mathcal{F}^\bullet)\otimes\check{H}^\bullet(Y;H^0\mathcal{G}^\bullet) \arrow[r] & \check{H}^\bullet(X\times Y;H^0\mathcal{H}^\bullet)
\end{tikzcd} \tag{4.8}$$

where the vertical arrows are the isomorphisms from [Par15, Proposition A.5.4] and the bottom horizontal arrow is the usual Künneth cup product map.

Proof. We consider the following commutative diagram:

$$\begin{tikzcd}
H^\bullet\mathcal{F}^\bullet(X)\otimes H^\bullet\mathcal{G}^\bullet(Y) \arrow[r] \arrow[d, "\sim"] & H^\bullet\left[\mathcal{F}^\bullet(X)\otimes\mathcal{G}^\bullet(Y)\right] \arrow[r] \arrow[d] & H^\bullet\mathcal{H}^\bullet(X\times Y) \arrow[d, "\sim"] \\
\check{H}^\bullet(X;\mathcal{F}^\bullet)\otimes\check{H}^\bullet(Y;\mathcal{G}^\bullet) \arrow[r] & H^\bullet\left[\check{C}^\bullet(X;\mathcal{F}^\bullet)\otimes\check{C}^\bullet(Y;\mathcal{G}^\bullet)\right] \arrow[r] & \check{H}^\bullet(X\times Y,\mathcal{H}^\bullet) \\
\check{H}^\bullet(X;\tau_{\le0}\mathcal{F}^\bullet)\otimes\check{H}^\bullet(Y;\tau_{\le0}\mathcal{G}^\bullet) \arrow[r] \arrow[u, "\sim"] \arrow[d, "\sim"] & H^\bullet\left[\check{C}^\bullet(X;\tau_{\le0}\mathcal{F}^\bullet)\otimes\check{C}^\bullet(Y;\tau_{\le0}\mathcal{G}^\bullet)\right] \arrow[r] \arrow[u] \arrow[d] & \check{H}^\bullet(X\times Y,\tau_{\le0}\mathcal{H}^\bullet) \arrow[u, "\sim"] \arrow[d, "\sim"] \\
\check{H}^\bullet(X;H^0\mathcal{F}^\bullet)\otimes\check{H}^\bullet(Y;H^0\mathcal{G}^\bullet) \arrow[r] & H^\bullet\left[\check{C}^\bullet(X;H^0\mathcal{F}^\bullet)\otimes\check{C}^\bullet(Y;H^0\mathcal{G}^\bullet)\right] \arrow[r] & \check{H}^\bullet(X\times Y,H^0\mathcal{H}^\bullet)
\end{tikzcd} \tag{4.9}$$

The far right and far left vertical maps are all isomorphisms by the same reasoning used in [Par15, Proposition A.5.4]. The outer square is exactly the desired diagram, so we are done.

Note that it is important to write $\mathcal{F}^\bullet\leftarrow\tau_{\le0}\mathcal{F}^\bullet\to H^0\mathcal{F}^\bullet$ and not $\mathcal{F}^\bullet\to\tau_{\ge0}\mathcal{F}^\bullet\leftarrow H^0\mathcal{F}^\bullet$, since there is a natural map $\tau_{\le0}\mathcal{F}^\bullet\otimes\tau_{\le0}\mathcal{G}^\bullet\to\tau_{\le0}\mathcal{H}^\bullet$, but no natural map $\tau_{\ge0}\mathcal{F}^\bullet\otimes\tau_{\ge0}\mathcal{G}^\bullet\to\tau_{\ge0}\mathcal{H}^\bullet$. ∎

Lemma 4.4. *Under the isomorphism* (4.3)*, the action of* (4.7) *induces the usual Künneth cup product map $\check{H}^\bullet(X;\mathfrak{o}_{X\,\mathrm{rel}\,\partial})\otimes\check{H}^\bullet(Y;\mathfrak{o}_{Y\,\mathrm{rel}\,\partial})\to\check{H}^\bullet(X\times Y;\mathfrak{o}_{X\times Y\,\mathrm{rel}\,\partial})$.*

Proof. By Lemma 4.3, this follows from [Par15, Definition 6.3.2]. ∎

Lemma 4.5. *Let $(X, \partial X) \to (\mathcal{S}, \partial\mathcal{S})$ (resp. $(Y, \partial Y) \to (\mathcal{T}, \partial\mathcal{T})$) be equipped with a finite locally orientable implicit atlas with boundary and cell-like stratification A (resp. B). Then the following diagram commutes:*

$$\begin{tikzcd} H^\bullet_{\mathrm{vir}}(X, \mathcal{S}; A) \otimes H^\bullet_{\mathrm{vir}}(Y, \mathcal{T}; B) \arrow[r] \arrow[d] & H^\bullet_{\mathrm{vir}}(X \times Y, \mathcal{S} \times \mathcal{T}; A \sqcup B) \arrow[d] \\ \check{H}^\bullet(X; \mathfrak{o}_X) \otimes \check{H}^\bullet(Y; \mathfrak{o}_Y) \arrow[r] & \check{H}^\bullet(X \times Y; \mathfrak{o}_{X\times Y}) \end{tikzcd} \tag{4.10}$$

Proof. By Lemma 4.3, this follows from [Par15, Proposition 6.2.3 and Definition 6.3.2]. □

4.2 Sketch of the construction

We first sketch the construction of virtual fundamental cycles on the moduli spaces $\overline{\mathcal{M}}_*(T)$. The remainder of this section is then devoted to turning this sketch into an actual proof.

The virtual fundamental cycle of a space with implicit atlas is represented (on the chain level) by the pushforward map (4.4). Since each space $\overline{\mathcal{M}}_*(T)$ is equipped with an implicit atlas $\bar{A}_*(T)$, the VFC machinery provides us with chain maps:

$$C^{\mathrm{vdim}(T)+\bullet}_{\mathrm{vir}}(\overline{\mathcal{M}}_*(T) \,\mathrm{rel}\, \partial; \bar{A}_*(T)) \to C_{-\bullet}(E; \bar{A}_*(T)) \tag{4.11}$$

Now the domain is quasi-isomorphic to Čech cochains $\check{C}^{\mathrm{vdim}(T)+\bullet}(\overline{\mathcal{M}}_*(T) \,\mathrm{rel}\, \partial; \mathfrak{o}_T)$ by (4.3) (recall $\mathfrak{o}_{\overline{\mathcal{M}}_*(T)} = \mathfrak{o}_T$), and the target is quasi-isomorphic to $\mathbb{Q}$. Thus (up to quasi-isomorphism), we have a collection of chain maps:

$$\check{C}^{\mathrm{vdim}(T)+\bullet}(\overline{\mathcal{M}}_*(T) \,\mathrm{rel}\, \partial; \mathfrak{o}_T) \to \mathbb{Q} \tag{4.12}$$

which is nothing other than a collection of (virtual fundamental) cycles:

$$[\overline{\mathcal{M}}_*(T)]^{\mathrm{vir}} \in \overline{C}_{\mathrm{vdim}(T)}(\overline{\mathcal{M}}_*(T) \,\mathrm{rel}\, \partial; \mathfrak{o}_T^\vee) \tag{4.13}$$

We remark that the dual of Čech cochains $\check{C}^\bullet$ is Steenrod chains $\overline{C}_\bullet$ (see [Par15, §A.9] and the references therein), although this level of precision is not needed in the present sketch.

Now supposing that (4.11) is sufficiently functorial with respect to morphisms and concatentations in $\mathcal{S}_*$, we conclude that the virtual fundamental cycles (4.13) satisfy the following identities:

$$\partial[\overline{\mathcal{M}}_*(T)]^{\mathrm{vir}} = \sum_{\mathrm{codim}(T'/T)=1} \frac{1}{|\mathrm{Aut}(T'/T)|}[\overline{\mathcal{M}}_*(T')]^{\mathrm{vir}} \tag{4.14}$$

$$[\overline{\mathcal{M}}_*(\#_i T_i)]^{\mathrm{vir}} = \frac{1}{|\mathrm{Aut}(\{T_i\}_i/\#T_i)|}\prod_i [\overline{\mathcal{M}}_*(T_i)]^{\mathrm{vir}} \tag{4.15}$$

which clearly imply the desired master equations. It also follows from the formal properties of the VFC package that $[\overline{\mathcal{M}}_*(T)]^{\mathrm{vir}}$ is the usual fundamental cycle if $\overline{\mathcal{M}}_*(T)$ is cut out transversally.

To turn this sketch into an actual construction, we must carry out the above arguments on the chain level, with all the necessary chain level functoriality with respect to morphisms and concatenations of trees. This is a non-trivial task, since many of the chain maps we are given go in the wrong direction (basically, we need to "invert quasi-isomorphisms"). The rest of this section is devoted to performing the necessary algebraic manipulations.

4.3 $(\mathcal{S}_{\mathrm{I}}, \mathcal{S}_{\mathrm{II}}, \mathcal{S}_{\mathrm{III}}, \mathcal{S}_{\mathrm{IV}})$-modules

We begin by introducing the notion of an $\mathcal{S}_*$-module, which formalizes the way in which the moduli spaces $\overline{\mathcal{M}}_*(T)$ fit together under morphisms and concatenations of trees (namely (2.3)–(2.4)). In fact, many (or even most) of the objects introduced and studied earlier in this paper are $\mathcal{S}_*$-modules, as we explain in the examples which follow the definition. This notion plays a key organizational role in what follows.

Definition 4.6. An object $T \in \mathcal{S}_*$ will be called *effective* iff $\overline{\mathcal{M}}_*(T) \neq \varnothing$. Note that (1) for any morphism $T \to T'$, if T is effective, then so is T', and (2) for any concatenation $\{T_i\}_i$, $\#_i T_i$ is effective iff every T_i is effective.

For the remainder of this section, we will use $\mathcal{S}_*$ to denote the full subcategory spanned by effective objects.

Definition 4.7. An $\mathcal{S}_{\mathrm{I}}$-*module* X_{I} valued in a symmetric monoidal category $\mathcal{C}^{\otimes}$ consists of the following data:

i. A functor $X_{\mathrm{I}} : \mathcal{S}_{\mathrm{I}} \to \mathcal{C}$.

ii. For every concatenation of $\{T_i\}_i$ in $\mathcal{S}_{\mathrm{I}}$, a morphism:

$$\bigotimes_i X_{\mathrm{I}}(T_i) \to X_{\mathrm{I}}(\#_i T_i) \tag{4.16}$$

such that the following diagrams commute:

$$\begin{tikzcd} \bigotimes_i X_{\mathrm{I}}(T_i) \arrow[r, "(4.16)"] \arrow[d] & X_{\mathrm{I}}(\#_i T_i) \arrow[d] \\ \bigotimes_i X_{\mathrm{I}}(T_i') \arrow[r, "(4.16)"] & X_{\mathrm{I}}(\#_i T_i') \end{tikzcd} \qquad \begin{tikzcd} & \bigotimes_i X_{\mathrm{I}}(\#_j T_{ij}) \arrow[dr, "(4.16)"] & \\ \bigotimes_{i,j} X_{\mathrm{I}}(T_{ij}) \arrow[ur, "\otimes_i (4.16)"] \arrow[rr, "(4.16)"] & & X_{\mathrm{I}}(\#_{ij} T_{ij}) \end{tikzcd} \tag{4.17}$$

for any morphism of concatenations $\{T_i\}_i \to \{T_i'\}_i$ and any composition of concatenations, respectively.

A morphism of $\mathcal{S}_{\mathrm{I}}$-modules is a natural transformation of functors compatible with (4.16).

Example 4.8. The functor $\overline{\mathcal{M}}_{\mathrm{I}}$ is an $\mathcal{S}_{\mathrm{I}}$-module (valued in the category of compact Hausdorff spaces, with the product symmetric monoidal structure).

Example 4.9. The functor $\bar{A}_{\mathrm{I}}$ is a (contravariant) $\mathcal{S}_{\mathrm{I}}$-module (valued in the category of sets, with the disjoint union symmetric monoidal structure).

Example 4.10. The functor $\mathfrak{o}$ is an $\mathcal{S}_{\mathrm{I}}$-module (valued in the category of orientation lines and isomorphisms, with the super tensor product symmetric monoidal structure).

Example 4.11. The functor $T \mapsto (\mathcal{S}_{\mathrm{I}})_{/T}$ is an $\mathcal{S}_{\mathrm{I}}$-module (valued in the category of categories, with the product symmetric monoidal structure).

Definition 4.12. An $\mathcal{S}_{\mathrm{II}}$-*module* X_{II} valued in $\mathcal{C}^{\otimes}$ consists of the following data:

i. An $\mathcal{S}_{\mathrm{I}}^{+}$-module X_{I}^{+} valued in $\mathcal{C}^{\otimes}$.
ii. An $\mathcal{S}_{\mathrm{I}}^{-}$-module X_{I}^{-} valued in $\mathcal{C}^{\otimes}$.
iii. A functor $X_{\mathrm{II}} : \mathcal{S}_{\mathrm{II}} \to \mathcal{C}$.
iv. For every concatenation $\{T_i\}_i$ in $\mathcal{S}_{\mathrm{II}}$, a morphism:

$$\bigotimes_i X_{\mathrm{I/II}}(T_i) \to X_{\mathrm{II}}(\#_i T_i) \tag{4.18}$$

such that the following diagrams commute:

$$\begin{tikzcd} \bigotimes_i X_{\mathrm{I/II}}(T_i) \arrow[r] \arrow[d] & X_{\mathrm{I/II}}(\#_i T_i) \arrow[d] \\ \bigotimes_i X_{\mathrm{I/II}}(T_i') \arrow[r] & X_{\mathrm{I/II}}(\#_i T_i') \end{tikzcd} \qquad \begin{tikzcd} & \bigotimes_i X_{\mathrm{I/II}}(\#_j T_{ij}) \arrow[dr] & \\ \bigotimes_{i,j} X_{\mathrm{I/II}}(T_{ij}) \arrow[ur] \arrow[rr] & & X_{\mathrm{II}}(\#_{ij} T_{ij}) \end{tikzcd} \tag{4.19}$$

for any morphism or composition of concatenations, respectively.

A morphism of $\mathcal{S}_{\mathrm{II}}$-modules consists of natural transformations of functors compatible with (4.16), (4.18).

Remark 4.13. It would perhaps be more proper to speak of "$(\mathcal{S}_{\mathrm{I}}^{\pm}, \mathcal{S}_{\mathrm{II}})$-modules $(X_{\mathrm{I}}^{\pm}, X_{\mathrm{II}})$", though such notation rapidly becomes unwieldy. It will usually be clear from context what $X_{\mathrm{I}}^{\pm}$ are once we have specified X_{II}, though at times we will specify the pair $(X_{\mathrm{I}}^{\pm}, X_{\mathrm{II}})$ for sake of clarity.

Example 4.14. Examples 4.8–4.11 all generalize: $(\overline{\mathcal{M}}_{\mathrm{I}}^{\pm}, \overline{\mathcal{M}}_{\mathrm{II}})$ is an $(\mathcal{S}_{\mathrm{I}}^{\pm}, \mathcal{S}_{\mathrm{II}})$-module, $(\bar{A}_{\mathrm{I}}^{\pm}, \bar{A}_{\mathrm{II}})$ is an $(\mathcal{S}_{\mathrm{I}}^{\pm}, \mathcal{S}_{\mathrm{II}})$-module, etc.

Definition 4.15. An $\mathcal{S}_{\mathrm{III}}$*-module* X_{III} valued in $\mathcal{C}^{\otimes}$ consists of the following data:
i. An $\mathcal{S}_{\mathrm{I}}^{+}$-module X_{I}^{+} valued in $\mathcal{C}^{\otimes}$.
ii. An $\mathcal{S}_{\mathrm{I}}^{-}$-module X_{I}^{-} valued in $\mathcal{C}^{\otimes}$.
iii. An $(\mathcal{S}_{\mathrm{I}}^{\pm}, \mathcal{S}_{\mathrm{II}}^{t=0})$-module $(X_{\mathrm{I}}^{\pm}, X_{\mathrm{II}}^{t=0})$ valued in $\mathcal{C}^{\otimes}$.
iv. An $(\mathcal{S}_{\mathrm{I}}^{\pm}, \mathcal{S}_{\mathrm{II}}^{t=1})$-module $(X_{\mathrm{I}}^{\pm}, X_{\mathrm{II}}^{t=1})$ valued in $\mathcal{C}^{\otimes}$.
v. A functor $X_{\mathrm{III}} : S_{\mathrm{III}} \to \mathcal{C}$.
vi. For every concatenation $\{T_i\}_i$ in $\mathcal{S}_{\mathrm{III}}$, a morphism:

$$\bigotimes_i X_{\mathrm{I/II/III}}(T_i) \to X_{\mathrm{III}}(\# T_i) \tag{4.20}$$

satisfying the natural compatibility conditions, as in Definition 4.12.

A morphism of $\mathcal{S}_{\mathrm{III}}$-modules consists of natural transformations of functors compatible with (4.16), (4.18), (4.20).

Definition 4.16. An $\mathcal{S}_{\mathrm{IV}}$*-module* X_{IV} valued in $\mathcal{C}^{\otimes}$ consists of the following data:
i. $\mathcal{S}_{\mathrm{I}}^{i}$-modules X_{I}^{i} valued in $\mathcal{C}^{\otimes}$ for $0 \le i \le 2$.
ii. $(\mathcal{S}_{\mathrm{I}}^{i,j}, \mathcal{S}_{\mathrm{II}}^{ij})$-modules $(X_{\mathrm{I}}^{i,j}, X_{\mathrm{II}}^{ij})$ valued in $\mathcal{C}^{\otimes}$ for $0 \le i < j \le 2$.

iii. A functor $X_{\mathrm{IV}} : \mathcal{S}_{\mathrm{IV}} \to \mathcal{C}$.
iv. For every concatenation $\{T_i\}_i$ in $\mathcal{S}_{\mathrm{IV}}$, a map:

$$\bigotimes_i X_{\mathrm{I/II/IV}}(T_i) \to X_{\mathrm{IV}}(\#_i T_i) \tag{4.21}$$

satisfying the natural compatibility conditions, as in Definition 4.12.

A morphism of $\mathcal{S}_{\mathrm{IV}}$-modules consists of natural transformations of functors compatible with (4.16), (4.18), (4.21).

4.4 Sketch of the construction (revisited)

We now revisit the sketch §4.2 using the language of $\mathcal{S}_*$-modules.

Recall that the virtual fundamental cycles of the moduli spaces $\overline{\mathcal{M}}_*(T)$ come packaged via the map (4.11). Now the key coherence properties of these cycles (4.14)–(4.15) shall be encoded in the fact that (4.11) is a map of $\mathcal{S}_*$-modules. Thus our first task is to define the $\mathcal{S}_*$-modules $C^{\bullet+\mathrm{vdim}}_{\mathrm{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)$ and $C_\bullet(E_*)$ along with the pushforward map:

$$C^{\bullet+\mathrm{vdim}}_{\mathrm{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial) \to C_{-\bullet}(E_*) \tag{4.22}$$

We then must relate $C^{\bullet+\mathrm{vdim}}_{\mathrm{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)$ to $\check{C}^{\bullet+\mathrm{vdim}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial; \mathfrak{o})$ (with an appropriate $\mathcal{S}_*$-module structure), and we must relate $C_\bullet(E_*)$ to $\mathbb{Q}$ (with the trivial $\mathcal{S}_*$-module structure). If all goes well, this will give virtual fundamental cycles (4.13) satisfying (4.14)–(4.15).

In reality, it is somewhat cumbersome to make sense out of $\check{C}^{\bullet+\mathrm{vdim}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial; \mathfrak{o})$ as an $\mathcal{S}_*$-module, so we will take a shortcut taking advantage of the fact that we are only interested in integrating the constant function 1 (and not more general cohomology classes) over the virtual fundamental cycles. We will define an $\mathcal{S}_*$-module $\mathbb{Q}[\mathcal{S}_*]$, which can be thought of as "the subcomplex of $\check{C}^{\bullet+\mathrm{vdim}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial; \mathfrak{o})$ generated by the characteristic functions of closed strata", and we will construct a corresponding map of $\mathcal{S}_*$-modules:

$$\mathbb{Q}[\mathcal{S}_*] \to C^{\bullet+\mathrm{vdim}}_{\mathrm{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial) \tag{4.23}$$

Now combining this map with (4.22) and our understanding of $C_\bullet(E_*) \cong \mathbb{Q}$, we obtain a map of $\mathcal{S}_*$-modules:

$$\mathbb{Q}[\mathcal{S}_*] \to \mathbb{Q} \tag{4.24}$$

Morally speaking, this map gives the values of integrating the characteristic functions of closed strata over the virtual fundamental cycle of $\overline{\mathcal{M}}_*(T)$. Practically speaking, from any map (4.24) it is straightforward to read off virtual moduli counts satisfying the master equations.

4.5 $\mathcal{S}_*$-modules $\mathbb{Q}$ and $\mathbb{Q}[\mathcal{S}_*]$

We now introduce the two basic $\mathcal{S}_*$-modules $\mathbb{Q}$ and $\mathbb{Q}[\mathcal{S}_*]$. We also make the elementary but crucial observation that a map of $\mathcal{S}_*$-modules $\mathbb{Q}[\mathcal{S}_*] \to \mathbb{Q}$ is nothing other than a collection of virtual moduli counts satisfying the relevant master equations.

Definition 4.17 ($\mathcal{S}_*$-module $\mathbb{Q}$)**.** Denote by $\mathbb{Q}$ the $\mathcal{S}_*$-module defined by $\mathbb{Q}(T) = \mathbb{Q}$ for all $T \in \mathcal{S}_*$, where the pushforward maps are the identity and the concatenation maps are multiplication.

We motivate the definition of $\mathbb{Q}[\mathcal{S}_*]$ as follows. It is not hard to check that $\check{C}^\bullet(\overline{\mathcal{M}}_*(T) \operatorname{rel} \partial; \mathfrak{o}_T)$ is quasi-isomorphic to the following total complex:

$$\check{C}^\bullet(\overline{\mathcal{M}}_*(T); \mathfrak{o}_T) \to \bigoplus_{\operatorname{codim}(T'/T)=1} \check{C}^{\bullet-1}(\overline{\mathcal{M}}_*(T')/\operatorname{Aut}(T'/T); \mathfrak{o}_{T'}) \to \bigoplus_{\operatorname{codim}(T''/T)=2} \check{C}^{\bullet-2}(\overline{\mathcal{M}}_*(T'')/\operatorname{Aut}(T''/T); \mathfrak{o}_{T''}) \to \cdots \tag{4.25}$$

This chain model is convenient because it is manifestly an $\mathcal{S}_*$-module: a morphism $T' \to T$ clearly induces a map of complexes (4.25) of degree $\operatorname{codim}(T'/T)$, while it is not so clear how to canonically define a corresponding map $\check{C}^\bullet(\overline{\mathcal{M}}_*(T') \operatorname{rel} \partial) \to \check{C}^{\bullet+\operatorname{codim}(T'/T)}(\overline{\mathcal{M}}_*(T) \operatorname{rel} \partial)$. The complex $\mathbb{Q}[\mathcal{S}_*](T)$ which we now define should be thought of as the subcomplex of (4.25) spanned by the constant sections $1 \in \check{C}^0(\overline{\mathcal{M}}_*(T')/\operatorname{Aut}(T'/T))$ for $T' \to T$.

Definition 4.18 ($\mathcal{S}_*$-module $\mathbb{Q}[\mathcal{S}_*]$)**.** For $T \in \mathcal{S}_*$, we define:

$$\mathbb{Q}[\mathcal{S}_*](T) := \bigoplus_{T' \in |(\mathcal{S}_*)_{/T}|} \mathfrak{o}_{T'} \tag{4.26}$$

The term corresponding to T' lies in degree $-\operatorname{vdim}(T')$ (cohomological grading), and the action of the differential on $\mathfrak{o}_{T'}$ is given by the sum over $T'' \in |(\mathcal{S}_*)_{/T'}|$ of the boundary map $\mathfrak{o}_{T'} \to \mathfrak{o}_{T''}$ multiplied by $\operatorname{Aut}(T''/T)/(\operatorname{Aut}(T'/T)\operatorname{Aut}(T''/T'))$.

A morphism $T' \to T$ induces a map:

$$\mathbb{Q}[\mathcal{S}_*](T') \to \mathbb{Q}[\mathcal{S}_*](T) \tag{4.27}$$

via pushforward under the functor $(\mathcal{S}_*)_{/T'} \to (\mathcal{S}_*)_{/T}$ and multiplication by $|\operatorname{Aut}(T''/T)/\operatorname{Aut}(T''/T')|$ on $\mathfrak{o}_{T''}$ for $T'' \in (\mathcal{S}_*)_{/T'}$.

A concatenation $\{T_i\}_i$ induces a map:

$$\bigotimes_i \mathbb{Q}[\mathcal{S}_*](T_i) \to \mathbb{Q}[\mathcal{S}_*](\#_i T_i) \tag{4.28}$$

by consideration of the isomorphism $(\mathcal{S}_*)_{/\#_i T_i} = \prod_i (\mathcal{S}_*)_{/T_i}$, covered by the tautological identifications $\mathfrak{o}_{\#_i T_i'} = \bigotimes_i \mathfrak{o}_{T_i'}$ multiplied by $|\operatorname{Aut}(\{T_i'\}_i/\#_i T_i')| = |\operatorname{Aut}(\{T_i\}_i/\#_i T_i)|$.

Thus $\mathbb{Q}[\mathcal{S}_*]$ is an $\mathcal{S}_*$-module (i.e. $\mathbb{Q}[\mathcal{S}_\mathrm{I}]$ is an $\mathcal{S}_\mathrm{I}$-module, $(\mathbb{Q}[\mathcal{S}_\mathrm{I}^\pm], \mathbb{Q}[\mathcal{S}_\mathrm{II}])$ is an $(\mathcal{S}_\mathrm{I}^\pm, \mathcal{S}_\mathrm{II})$-module, etc.).

The factors appearing in the differential and structure maps above are explained by regarding $\mathbb{Q}[\mathcal{S}_*](T)$ as generated by elements called "$\mathbf{1}_{\overline{\mathcal{M}}(T')/\operatorname{Aut}(T'/T)} = |\operatorname{Aut}(T'/T)|^{-1}\,\mathbf{1}_{\overline{\mathcal{M}}(T')}$".

Remark 4.19. It is easy to check that there is a natural bijection between morphisms of $\mathcal{S}_*$-modules $\mathbb{Q}[\mathcal{S}_*] \to \mathbb{Q}$ and collections of virtual moduli counts $[\overline{\mathcal{M}}_*(T)]^{\mathrm{vir}} \in (\mathfrak{o}_T^\vee)^{\operatorname{Aut}(T)}$ for

$\mathrm{vdim}(T) = 0$ satisfying:

$$0 = \sum_{\mathrm{codim}(T'/T)=1} \frac{1}{|\mathrm{Aut}(T'/T)|} [\overline{\mathcal{M}}_*(T')]^{\mathrm{vir}} \tag{4.29}$$

$$[\overline{\mathcal{M}}_*(\#_i T_i)]^{\mathrm{vir}} = \frac{1}{|\mathrm{Aut}(\{T_i\}_i/\#T_i)|} \prod_i [\overline{\mathcal{M}}_*(T_i)]^{\mathrm{vir}} \tag{4.30}$$

(where $\overline{\mathcal{M}}_*(T)$ is interpreted as zero if $\mathrm{vdim}(T) \neq 0$). Note that $\mathrm{Aut}(T)$-invariance forces $[\overline{\mathcal{M}}_*(T)]^{\mathrm{vir}}$ to vanish if any of the input/output edges of T is labeled with a bad Reeb orbit.

A collection of such virtual moduli counts clearly gives virtual moduli counts as in Theorems I, II, III, IV satisfying the relevant master equation (1.4)/(1.13)/(1.19)/(1.26). Note that for $* = \mathrm{I}$ (resp. $* = \mathrm{III}$, $* = \mathrm{IV}$), this involves choosing an orientation on $\mathbb{R}$ (resp. $[0, 1]$, $[0, \infty]$).

4.6 $\mathcal{S}_*$-modules $C_{\mathrm{vir}}^{\bullet+\mathrm{vdim}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)$ and $C_\bullet(E_*)$

We now introduce the two $\mathcal{S}_*$-modules central to our definition of the virtual fundamental cycles, namely $C_{\mathrm{vir}}^{\bullet+\mathrm{vdim}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)$ and $C_\bullet(E_*)$.

Defining $C_{\mathrm{vir}}^{\bullet+\mathrm{vdim}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)$ as an $\mathcal{S}_*$-module is non-trivial for the following reason. We would like to associate to $T \in \mathcal{S}_*$ the complex $C_{\mathrm{vir}}^{\bullet+\mathrm{vdim}(T)}(\overline{\mathcal{M}}_*(T) \operatorname{rel} \partial; \bar{A}_*(T))$. However, a map $T' \to T$ does not induce a map on such complexes as desired, rather only a diagram:

$$\begin{tikzcd} C_{\mathrm{vir}}^{\bullet+\mathrm{vdim}(T')}(\overline{\mathcal{M}}_*(T') \operatorname{rel} \partial; \bar{A}_*(T)) \arrow[d, "\sim"] \arrow[dr] & \\ C_{\mathrm{vir}}^{\bullet+\mathrm{vdim}(T')}(\overline{\mathcal{M}}_*(T') \operatorname{rel} \partial; \bar{A}_*(T')) & C_{\mathrm{vir}}^{\bullet+\mathrm{vdim}(T)}(\overline{\mathcal{M}}_*(T) \operatorname{rel} \partial; \bar{A}_*(T)) \end{tikzcd} \tag{4.31}$$

Fortunately, this diagram also suggests a solution. Namely, we instead associate to $T \in \mathcal{S}_*$ a suitable homotopy colimit of $C_{\mathrm{vir}}^{\mathrm{vdim}(T')+\bullet}(\overline{\mathcal{M}}_*(T') \operatorname{rel} \partial, \bar{A}_*(T'))$ over $T' \in (\mathcal{S}_*)_{/T}$ which is quasi-isomorphic to $C_{\mathrm{vir}}^{\mathrm{vdim}(T)+\bullet}(\overline{\mathcal{M}}_*(T) \operatorname{rel} \partial; \bar{A}_*(T))$ and has natural pushforward maps for morphisms $T' \to T$.

Definition 4.20 (Homotopy diagram)**.** Let $\mathcal{S}$ be a finite category (meaning $\#|\mathcal{S}| < \infty$ and $\#\mathrm{Hom}(T_1, T_2) < \infty$ for $T_1, T_2 \in \mathcal{S}$). A *homotopy diagram* over $\mathcal{S}$ shall mean the following:

i. Let Δ^p denote the *simplex category* (so that a functor $\Delta^p \to \mathcal{S}$ is a chain of morphisms $T_0 \to \cdots \to T_p$). For every functor $\sigma : \Delta^p \to \mathcal{S}$, we specify a complex $A^\bullet(\sigma)$, and for every map of simplices $r : \Delta^q \to \Delta^p$, we specify a map $A^\bullet(\sigma) \to A^\bullet(\sigma \circ r)$. In other words, $A^\bullet$ is a functor from the category whose objects are morphisms $\sigma^p : \Delta^p \to \mathcal{S}$ and whose morphisms $\sigma^p \to \sigma^q$ are factorizations $\Delta^q \to \Delta^p \to \mathcal{S}$.

ii. We require that for any surjection $r : \Delta^q \twoheadrightarrow \Delta^p$, the induced map $A^\bullet(\sigma) \to A^\bullet(\sigma \circ r)$ is an isomorphism.

The coefficient systems $A^\bullet(T_0 \to \cdots \to T_p)$ relevant for us will actually only depend on the total composition $T_0 \to T_p$.

In the case that $\mathcal{S}$ is a poset and the homotopy diagram only depends on the total composition $T_0 \to T_p$, this notion of a homotopy diagram reduces to [Par15, Definition A.7.1].

Definition 4.21 (Homotopy colimit)**.** Let $\mathcal{S}$ be a finite category, and let $A^\bullet$ be a homotopy diagram over $\mathcal{S}$. We define:

$$\operatorname*{hocolim}_{\mathcal{S}} A^\bullet := \bigoplus_{p\geq 0} \bigoplus_{T_0\to\cdots\to T_p} A^\bullet(T_0 \to \cdots \to T_p)_{\operatorname{Aut}(T_0\to\cdots\to T_p)} \otimes \mathfrak{o}_{\Delta^p}[p] \tag{4.32}$$

where $T_0 \to \cdots \to T_p$ denotes a chain of *nontrivial* morphisms in $\mathcal{S}$ (i.e. morphisms which are not isomorphisms). This homotopy colimit is just chains on the nerve of $\mathcal{S}$, using $A^\bullet$ as the coefficient system (this interpretation of (4.32) also specifies the differential and explains the appearance of $\mathfrak{o}_{\Delta^p} \cong \mathbb{Z}$, the orientation line of the p-simplex).

In the case that $\mathcal{S}$ is a poset and the homotopy diagram only depends on the total composition $T_0 \to T_p$, this reduces to [Par15, Definition A.7.2].

Let $A^\bullet$ be a homotopy diagram over $\mathcal{S}$. Then a functor $f : \mathcal{T} \to \mathcal{S}$ induces a natural map:

$$\operatorname*{hocolim}_{\mathcal{T}} f^* A^\bullet \to \operatorname*{hocolim}_{\mathcal{S}} A^\bullet \tag{4.33}$$

Let $A_i^\bullet$ be a finite collection of homotopy diagrams over $\mathcal{S}_i$. Then there is an Eilenberg–Zilber quasi-isomorphism:

$$\bigotimes_i \operatorname*{hocolim}_{\mathcal{S}_i} A_i^\bullet \xrightarrow{\sim} \operatorname*{hocolim}_{\prod_i \mathcal{S}_i} \bigotimes_i A_i^\bullet \tag{4.34}$$

corresponding to the standard simplicial subdivision of $\Delta^p \times \Delta^q$ into $\binom{p+q}{p}$ copies of Δ^{p+q}.

Definition 4.22 ($\mathcal{S}_*$-module $C^{\bullet+\operatorname{vdim}}_{\operatorname{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)$)**.** For $T \in \mathcal{S}_*$, we define:

$$C^\bullet_{\operatorname{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)(T) := \operatorname*{hocolim}_{T''\to T' \in (\mathcal{S}_*)_{/T}} C^{\bullet - \operatorname{codim}(T''/T)}_{\operatorname{vir}}(\overline{\mathcal{M}}_*(T'') \operatorname{rel} \partial, \bar{A}_*(T')) \tag{4.35}$$

This is the homotopy colimit over the category $(\mathcal{S}_*)_{/T}$, with a coefficient system depending only on the total composition $T_0 \to T_p$. The structure maps of the coefficient system are the maps illustrated in (4.31). The natural inclusion map:

$$C^\bullet_{\operatorname{vir}}(\overline{\mathcal{M}}_*(T) \operatorname{rel} \partial; \bar{A}_*(T)) \hookrightarrow C^\bullet_{\operatorname{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)(T) \tag{4.36}$$

is a quasi-isomorphism by [Par15, Lemma A.7.3], which uses the fact that each leftmost map in (4.31) is a quasi-isomorphism. Thus this homotopy colimit recognizes the fact that $T \in (\mathcal{S}_*)_{/T}$ is a final object.

Now $C^{\bullet+\operatorname{vdim}}_{\operatorname{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)$ naturally has the structure of an $\mathcal{S}_*$-module, as follows. A morphism $T \to T'$ induces a natural pushforward map (4.33) induced by the functor $(\mathcal{S}_*)_{/T} \to (\mathcal{S}_*)_{/T'}$, which is covered by a natural isomorphism of homotopy diagrams. Given a concatenation $\{T_i\}_i$, there is a natural Eilenberg–Zilber map (4.34) induced by the isomorphism $(\mathcal{S}_*)_{/\#_i T_i} = \prod_i (\mathcal{S}_*)_{/T_i}$, which is covered by a morphism of homotopy diagrams coming from the product maps (4.7).

Definition 4.23 ($\mathcal{S}_*$-module $C_\bullet(E_*)$)**.** For $T \in \mathcal{S}_*$, we define:

$$C_\bullet(E_*)(T) := \operatorname*{hocolim}_{T' \in (\mathcal{S}_*)_{/T}} C_\bullet(E; \bar{A}_*(T')) \tag{4.37}$$

This is the homotopy colimit over the category $(\mathcal{S}_*)_{/T}$, with a coefficient system depending only on the last object T_p. Again by [Par15, Lemma A.7.3], we have:

$$H_\bullet(E_*)(T) = H_\bullet(E; \bar{A}_*(T)) = \mathbb{Q} \tag{4.38}$$

As in Definition 4.22, $C_\bullet(E_*)$ naturally has the structure of a $\mathcal{S}_*$-module. The isomorphism above is in fact an isomorphism of $\mathcal{S}_*$-modules $H_\bullet(E_*) = \mathbb{Q}$.

There is a canonical map of $\mathcal{S}_*$-modules:

$$C_{\mathrm{vir}}^{\bullet+\mathrm{vdim}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial) \to C_{-\bullet}(E_*) \tag{4.39}$$

induced by (4.4).

Definition 4.24 ($\mathcal{S}_*$-module $\operatorname{Hom}_{(\mathcal{S}_*)_{/T}}(\mathbb{Q}[\mathcal{S}_*], C_{\mathrm{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial))$)**.** For $T \in \mathcal{S}_*$, we consider:

$$\operatorname{Hom}_{(\mathcal{S}_*)_{/T}}(\mathbb{Q}[\mathcal{S}_*], C_{\mathrm{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)) \tag{4.40}$$

Namely, an element of this group is a natural transformation of functors from $(\mathcal{S}_*)_{/T}$ to the category of *graded $\mathbb{Q}$-vector spaces*, and we equip it with the usual differential $d \circ f - (-1)^{|f|} f \circ d$.

Now (4.40) is a contravariant $\mathcal{S}_*$-module as follows. A map $T \to T'$ clearly gives a restriction map from homomorphisms over $(\mathcal{S}_*)_{/T'}$ to homomorphisms over $(\mathcal{S}_*)_{/T}$. A concatenation $\{T_i\}_i$ induces a map in the correct direction by virtue of the fact that the concatenation maps for $\mathbb{Q}[\mathcal{S}_*]$ are isomorphisms.

Lemma 4.25. *The cohomology of* (4.40) *is canonically isomorphic as an $\mathcal{S}_*$-module to $T \mapsto \check{H}^\bullet(\overline{\mathcal{M}}_{\mathrm{I}}(T))$ equipped with the Künneth product multiplied by $|\operatorname{Aut}(\{T_i\}_i/\#_i T_i)|$.*

Proof. First, observe that:

$$\operatorname{Hom}_{(\mathcal{S}_*)_{/T}}(\mathbb{Q}[\mathcal{S}_*], C_{\mathrm{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial)) = \prod_{T' \to T} \left[\mathfrak{o}_{T'}^\vee \otimes C_{\mathrm{vir}}^\bullet(\overline{\mathcal{M}}_{\mathrm{I}} \operatorname{rel} \partial)(T')\right]^{\operatorname{Aut}(T'/T)} \tag{4.41}$$

Now there are natural quasi-isomorphisms:

$$\prod_{T' \to T} \left[\mathfrak{o}_{T'}^\vee \otimes \operatorname*{hocolim}_{T''' \to T'' \in (\mathcal{S}_*)_{/T'}} C_{\mathrm{vir}}^{\bullet - \operatorname{codim}(T'''/T')}(\overline{\mathcal{M}}_{\mathrm{I}}(T''') \operatorname{rel} \partial; \bar{A}_{\mathrm{I}}(T''))\right]^{\operatorname{Aut}(T'/T)} \tag{4.42}$$

$$\uparrow$$

$$\prod_{T' \to T} \left[\mathfrak{o}_{T'}^\vee \otimes \operatorname*{hocolim}_{T''' \to T'' \in (\mathcal{S}_*)_{/T'}} C_{\mathrm{vir}}^{\bullet - \operatorname{codim}(T'''/T')}(\overline{\mathcal{M}}_{\mathrm{I}}(T''') \operatorname{rel} \partial; \bar{A}_{\mathrm{I}}(T))\right]^{\operatorname{Aut}(T'/T)} \tag{4.43}$$

$$\downarrow$$

$$\prod_{T' \to T} \left[\mathfrak{o}_{T'}^\vee \otimes C_{\mathrm{vir}}^\bullet(\overline{\mathcal{M}}_{\mathrm{I}}(T') \operatorname{rel} \partial; \bar{A}_{\mathrm{I}}(T))\right]^{\operatorname{Aut}(T'/T)} \tag{4.44}$$

The first map is increasing the atlas via (4.5), and it is clearly a filtered quasi-isomorphism. The second map collapses the hocolim (i.e. the natural pushforward on $p = 0$ direct summands and zero for $p > 0$); it is a quasi-isomorphism because each inclusion (4.36) is a quasi-isomorphism.

Now [Par15, Proposition 6.2.3] implies that (4.44) is canonically isomorphic to $\check{H}^\bullet(\overline{\mathcal{M}}_{\mathrm{I}}(T))$ as a functor on $\mathcal{S}_*$, and Lemma 4.5 implies that the concatenation maps agree with the Künneth cup product maps. Note that we are applying these results in slightly greater generality than they are stated, since the stratifications in question are merely locally cell-like and not globally cell-like; nevertheless this is not an issue since their proofs are entirely local. $\square$

Lemma 4.26. *There exists a map of $\mathcal{S}_*$-modules:*

$$\mathbb{Q}[\mathcal{S}_*] \to C^{\bullet+\mathrm{vdim}}_{\mathrm{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial) \tag{4.45}$$

with the following property. Such a map determines a cycle in (4.40) for all $T \in \mathcal{S}_*$, and thus by Lemma 4.25 determines an element of $\check{H}^0(\overline{\mathcal{M}}_*(T))$. We require that this coincide with the class of the constant function $1 \in \check{H}^0(\overline{\mathcal{M}}_*(T))$ for all $T \in \mathcal{S}_*$.

Proof. We argue by induction on T. For T non-maximal, write $T = \#_i T_i$ with T_i maximal. For non-maximal T, cofibrancy of $\mathbb{Q}[\mathcal{S}_*]$ (see Lemma 4.29) determines the map completely on $\mathbb{Q}[\mathcal{S}_*](T)$, and the map has the desired property by Lemma 4.25. For maximal T, consider the following restriction map:

$$\prod_{T'\to T} \left[\mathfrak{o}^\vee_{T'} \otimes C^\bullet_{\mathrm{vir}}(\overline{\mathcal{M}}_{\mathrm{I}} \operatorname{rel} \partial)(T')\right]^{\mathrm{Aut}(T'/T)} \to \prod_{\mathrm{codim}(T'/T)\geq 1} \left[\mathfrak{o}^\vee_{T'} \otimes C^\bullet_{\mathrm{vir}}(\overline{\mathcal{M}}_{\mathrm{I}} \operatorname{rel} \partial)(T')\right]^{\mathrm{Aut}(T'/T)} \tag{4.46}$$

The part of the map defined thus far determines a cycle on the right, which we must lift to a cycle on the left in a particular cohomology class. Since (4.46) is surjective, it suffices to show that the desired cohomology class on the left is mapped to the given cohomology class on the right. Now the argument in the proof of Lemma 4.25 shows that on cohomology, the map (4.46) is just the restriction map $\check{H}^\bullet(\overline{\mathcal{M}}_{\mathrm{I}}(T)) \to \check{H}^\bullet(\partial\overline{\mathcal{M}}_{\mathrm{I}}(T))$. So by the sheaf property for $\check{H}^0$, it suffices to show equality after further restriction to $\check{H}^\bullet(\overline{\mathcal{M}}_{\mathrm{I}}(T'))$ for all nontrivial morphisms $T' \to T$. This holds by the induction hypothesis, so we are done. $\square$

4.7 Cofibrant $\mathcal{S}_*$-modules

We would like to upgrade the isomorphism of $\mathcal{S}_*$-modules $H_\bullet(E_*) = \mathbb{Q}$ to a quasi-isomorphism of $\mathcal{S}_*$-modules $C_\bullet(E_*) \xrightarrow{\sim} \mathbb{Q}$. Unfortunately, the natural strategy to construct such a quasi-isomorphism, namely by induction on T, fails. As a substitute, we introduce the notion of a *cofibrant* $\mathcal{S}_*$-module, and we construct a cofibrant $\mathcal{S}_*$-module $C^{\mathrm{cof}}_\bullet(E_*)$ with a natural quasi-isomorphism $C^{\mathrm{cof}}_\bullet(E_*) \xrightarrow{\sim} C_\bullet(E_*)$. Now cofibrancy is exactly the condition which we need to see through the inductive construction of a quasi-isomorphism $C^{\mathrm{cof}}_\bullet(E_*) \xrightarrow{\sim} \mathbb{Q}$. The resulting diagram:

$$C_\bullet(E_*) \xleftarrow{\sim} C^{\mathrm{cof}}_\bullet(E_*) \xrightarrow{\sim} \mathbb{Q} \tag{4.47}$$

turns out to be enough for our purposes.

Definition 4.27 (Cofibrant $\mathcal{S}_*$-module)**.** An $\mathcal{S}_*$-module X_* shall be called *cofibrant* iff it satisfies the following two properties:

i. For all concatenations $\{T_i\}_i$ in $\mathcal{S}_*$, the induced map:

$$\Big[\bigotimes_i X_*(T_i)\Big]_{\mathrm{Aut}(\{T_i\}_i/\#T_i)} \xrightarrow{\sim} X_*(\#_i T_i) \tag{4.48}$$

is an isomorphism. Note that this follows if (4.16)/(4.18)/(4.20)/(4.21) is itself an isomorphism.

ii. For maximal $T \in \mathcal{S}_*$, the map:

$$\operatorname*{colim}_{\mathrm{codim}(T'/T)\geq 1} X_*(T') \rightarrowtail X_*(T) \tag{4.49}$$

is injective. More precisely, the left side denotes the colimit over the full subcategory of $(\mathcal{S}_*)_{/T}$ spanned by objects $T' \to T$ with $\mathrm{codim}(T'/T) \geq 1$. In fact, injectivity of (4.49) for maximal T implies injectivity for all T, as we now argue. Indeed, fix $T \in \mathcal{S}_*$, and write $T = \#_i T_i$ for maximal trees T_i. Now we have $(\mathcal{S}_*)_{/T} = \prod_i (\mathcal{S}_*)_{/T_i}$. Consider the cubical diagram:

$$\bigotimes_i \Big[\operatorname*{colim}_{\mathrm{codim}(T_i'/T_i)\geq 1} X_*(T_i') \rightarrowtail X_*(T_i)\Big] \tag{4.50}$$

Now (4.49) for T is precisely the map to the top vertex of the cube (4.50) from the colimit over its remaining vertices. This map is clearly injective given that each map in (4.50) is injective.

Note that for $(\mathcal{S}_{\mathrm{I}}^{\pm}, \mathcal{S}_{\mathrm{II}})$-modules, the above conditions are imposed over each of $\mathcal{S}_{\mathrm{I}}^{\pm}$, $\mathcal{S}_{\mathrm{II}}$, and so on.

Cofibrancy of an $\mathcal{S}_*$-module X_* is important because it allows us to construct maps out of X_* by induction on $T \in \mathcal{S}_*$, partially ordered as in Definition 4.30.

Remark 4.28. If we were working over $\mathbb{Z}$, it would be important to require that (4.49) be injective *with projective cokernel*.

Lemma 4.29. *The $\mathcal{S}_*$-module $\mathbb{Q}[\mathcal{S}_*]$ is cofibrant.*

Proof. The concatenation maps are isomorphisms by definition since $(\mathcal{S}_*)_{/\#_i T_i} = \prod_i (\mathcal{S}_*)_{/T_i}$. Let us now show that:

$$\operatorname*{colim}_{\mathrm{codim}(T'/T)\geq 1} \mathbb{Q}[\mathcal{S}_*](T') \rightarrowtail \mathbb{Q}[\mathcal{S}_*](T) \tag{4.51}$$

is an isomorphism onto the subspace of $\mathbb{Q}[\mathcal{S}_*](T)$ generated by those $T' \in |(\mathcal{S}_*)_{/T}|$ of codimension ≥ 1 (certainly this is sufficient).

Both sides of (4.51) are graded by $|(\mathcal{S}_*)_{/T}|$, so it suffices to fix $T' \in |(\mathcal{S}_*)_{/T}|$ of codimension ≥ 1 and show that (4.51) is an isomorphism on T'-graded pieces. The map (4.51) is certainly surjective onto the T'-graded piece $\mathfrak{o}_{T'} \subseteq \mathbb{Q}[\mathcal{S}_*](T)$. Moreover, there is a section from $\mathfrak{o}_{T'}$ back to the left side of (4.51) via the inclusion $\mathfrak{o}_{T'} \subseteq \mathbb{Q}[\mathcal{S}_*](T')$. Now it is enough to argue that this section is surjective (onto the T'-graded piece), but this is clear since $\mathbb{Q}[\mathcal{S}_*](T'')$ contributes to the T'-graded piece exactly when there is a factorization $T' \to T'' \to T$. ∎

Definition 4.30 (Partial order on $|\mathcal{S}_*|$)**.** For $T, T' \in \mathcal{S}_*$, let us write $T' \preceq T$ iff there is a morphism $\#_i T_i \to T$ with some T_i isomorphic to T'. We claim that $\preceq$ is a partial order. Indeed, reflexivity and transitivity are immediate. Antisymmetry follows from our restriction to effective trees and compactness. Compactness also implies that the partial order $\preceq$ is *well-founded* (i.e. there is no infinite strictly decreasing sequence $T_1 \succ T_2 \succ \cdots$), and hence induction on $|\mathcal{S}_*|$ partially ordered by $\preceq$ is justified.

Definition 4.31 (Cofibrant $\mathcal{S}_*$-module $C_\bullet^{\mathrm{cof}}(E_*)$)**.** We now define a cofibrant $\mathcal{S}_*$-module $C_\bullet^{\mathrm{cof}}(E_*)$ together with a quasi-isomorphism:

$$C_\bullet^{\mathrm{cof}}(E_*) \to C_\bullet(E_*) \tag{4.52}$$

which is surjective for maximal T. Furthermore, the action of the "paths between basepoints" subgroup of $\operatorname{Aut}(T)$ on $C_\bullet^{\mathrm{cof}}(E_*)(T)$ will be trivial for all T (as it is for $C_\bullet(E_*)(T)$).

We construct $C_\bullet^{\mathrm{cof}}(E_*)(T)$ by induction on T, partially ordered as in Definition 4.30. For T non-maximal, write $T = \#_i T_i$ with T_i maximal. Then the definition of cofibrancy both forces us to take $C_\bullet^{\mathrm{cof}}(E_*)(T) := \bigotimes_i C_\bullet^{\mathrm{cof}}(E_*)(T_i)$ and assures that the $\mathcal{S}_*$-module structure maps with target $C_\bullet^{\mathrm{cof}}(E_*)(T)$ exist and are unique. For T maximal, consider the following diagram:

$$\begin{tikzcd}
\operatorname*{colim}_{\operatorname{codim}(T'/T)\geq 1} C_\bullet^{\mathrm{cof}}(E_*)(T') \arrow[r, tail, dashed] \arrow[d, "(4.52)"] & C_\bullet^{\mathrm{cof}}(E_*)(T) \arrow[d, two heads, dashed, "(4.52)", "\sim"'] \\
\operatorname*{colim}_{\operatorname{codim}(T'/T)\geq 1} C_\bullet(E_*)(T') \arrow[r] & C_\bullet(E_*)(T)
\end{tikzcd} \tag{4.53}$$

We define $C_\bullet^{\mathrm{cof}}(E_*)(T)$ to be the mapping cylinder of the composition of the two solid maps, which clearly fits into the diagram as desired. Now the top horizontal map defines the $\mathcal{S}_*$-module structure maps with target $C_\bullet^{\mathrm{cof}}(E_*)(T)$, and the commutativity of the diagram ensures that (4.52) is a map of $\mathcal{S}_*$-modules.

Lemma 4.32. *Let $\mathcal{T} \subseteq (\mathcal{S}_*)_{/T_a}$ be a full subcategory which is "downward closed" in the sense that $T' \to T$ and $T \in \mathcal{T}$ implies $T' \in \mathcal{T}$. Then the natural map:*

$$C_\bullet N_\bullet \mathcal{T} \cong \operatorname*{hocolim}_{T\in\mathcal{T}} C_\bullet^{\mathrm{cof}}(E_*)(T) \to \operatorname*{colim}_{T\in\mathcal{T}} C_\bullet^{\mathrm{cof}}(E_*)(T) \tag{4.54}$$

is a quasi-isomorphism. In particular, the homology of the right side is supported in degrees ≥ 0, and in degree zero is generated by the images of $H_0^{\mathrm{cof}}(E_*)(T)$ for $T \in \mathcal{T}$.

Proof. We argue by induction on $\#\mathcal{T}$ (note that $|(\mathcal{S}_*)_{/T}|$ is finite), the case $\#\mathcal{T} = 0$ being trivial. Pick a maximal object $T_0 \in \mathcal{T}$, so that (abbreviating $X := C_\bullet^{\mathrm{cof}}(E_{\mathrm{I}})$) there is a short exact sequence:

$$0 \to \operatorname*{colim}_{T\in\mathcal{T}^{<T_0}} X(T) \to X(T_0)_{\operatorname{Aut}(T_0)} \oplus \operatorname*{colim}_{T\in\mathcal{T}\setminus\{T_0\}} X(T) \to \operatorname*{colim}_{T\in\mathcal{T}} X(T) \to 0 \tag{4.55}$$

Here $\operatorname{Aut}(T_0)$ denotes the automorphism group of T_0 as an object of $\mathcal{T}$. Right exactness is clear, and left exactness holds by injectivity of (4.49) for T_0 and exactness of coinvariants

$_{\mathrm{Aut}(T_0)}$ in the category of $\mathbb{Q}$-vector spaces. Now there is an associated morphism of long exact sequences from the long exact sequence of hocolim to the long exact sequence of colim induced by (4.55). Now the induction hypothesis in combination with the five lemma gives the desired result. ☐

Lemma 4.33. *There exists a quasi-isomorphism of $\mathcal{S}_*$-modules:*

$$C_\bullet^{\mathrm{cof}}(E_*) \xrightarrow{\sim} \mathbb{Q} \tag{4.56}$$

inducing the canonical isomorphism $H_\bullet^{\mathrm{cof}}(E_*) = H_\bullet(E_*) = \mathbb{Q}$ from Definition 4.23.

Proof. We argue by induction on $T \in \mathcal{S}_*$, partially ordered as in Definition 4.30. For T non-maximal, write $T = \#_i T_i$ for T_i maximal. Then cofibrancy of $C_\bullet^{\mathrm{cof}}(E_*)$ both forces us to take $p_*(T) := \bigotimes_i p_*(T_i)$ and assures that this choice is compatible with the maps defined thus far. For maximal T, we would like to fill in the diagram:

$$\begin{tikzcd} \operatorname*{colim}_{\mathrm{codim}(T'/T)\geq 1} C_\bullet^{\mathrm{cof}}(E_*)(T') \arrow[r, tail] \arrow[d, "p_*"] & C_\bullet^{\mathrm{cof}}(E_*)(T) \arrow[dl, dashed, "p_*"] \\ \mathbb{Q} & \end{tikzcd} \tag{4.57}$$

with a map $C_\bullet^{\mathrm{cof}}(E_*)(T) \to \mathbb{Q}$ in a particular chain homotopy class. The horizontal map is injective since $C_\bullet^{\mathrm{cof}}(E_*)$ is cofibrant; it follows that it is enough to show that the diagram commutes up to chain homotopy. Let $C_\bullet$ stand for the colimit above; then the map $H_0 \operatorname{Hom}(C_\bullet, \mathbb{Q}) \to \operatorname{Hom}(H_0 C_\bullet, \mathbb{Q})$ is an isomorphism.[14] Thus it suffices to show that (4.57) commutes on homology, which follows from Lemma 4.32. Finally, we may ensure p_* is $\mathrm{Aut}(T)$-invariant by averaging (this is necessary for p_* to be a natural transformation of functors). ☐

4.8 Sets Θ_{I}, Θ_{II}, Θ_{III}, Θ_{IV}

We now conclude by defining the sets Θ_{I}, Θ_{II}, Θ_{III}, Θ_{IV} which index all possible choices of the extra data necessary to fix coherent virtual fundamental cycles on the moduli spaces $\overline{\mathcal{M}}_{\mathrm{I}}$, $\overline{\mathcal{M}}_{\mathrm{II}}$, $\overline{\mathcal{M}}_{\mathrm{III}}$, $\overline{\mathcal{M}}_{\mathrm{IV}}$.

Definition 4.34 (Sets Θ_*)**.** An element $\theta \in \Theta_*$ consists of two pieces of data.

The first piece of data is a choice of finite $\mathrm{Aut}(T)$-invariant subatlases $B_*(T) \subseteq A_*(T)$ on $\overline{\mathcal{M}}_*(T)$ for all maximal $T \in \mathcal{S}_*$ (e.g. if $* = \mathrm{II}$, we choose $B_{\mathrm{I}}^{\pm}(T) \subseteq A_{\mathrm{I}}^{\pm}(T)$ and $B_{\mathrm{II}}(T) \subseteq A_{\mathrm{II}}(T)$). Now we define $\bar{B}_*(T)$ as in (3.1), and we use $\bar{B}_*$ in place of $\bar{A}_*$ in the definitions from §§4.6–4.7.

[14] If we were working over $\mathbb{Z}$, this map would be surjective with kernel $\operatorname{Ext}^1(H_{-1}C_\bullet, \mathbb{Z})$, and we would need to invoke Lemma 4.32 to see that $H_{-1}C_\bullet$ vanishes.

The second piece of data is a commuting diagram[15] of $\mathcal{S}_*$-modules:

$$\begin{tikzcd} \mathbb{Q}[\mathcal{S}_*] \arrow[rr, "\tilde{w}_*"] \arrow[d, "w_*"] & & C^{\mathrm{cof}}_{-\bullet}(E_*) \arrow[r, "p_*"] \arrow[d, two heads, "\sim"', "(4.52)"] & \mathbb{Q} \\ C^{\bullet+\mathrm{vdim}}_{\mathrm{vir}}(\overline{\mathcal{M}}_* \operatorname{rel} \partial) \arrow[rr, "(4.39)"] & & C_{-\bullet}(E_*) & \end{tikzcd} \tag{4.58}$$

satisfying the following properties:

- We require that p_* induce the canonical isomorphism $H_\bullet(E_*) = H^{\mathrm{cof}}_\bullet(E_*) = \mathbb{Q}$ from Definition 4.23.
- We require that w_* satisfy the conclusion of Lemma 4.26.

Note that there are natural forgetful morphisms:

$$\Theta_{\mathrm{II}} \to \Theta^+_{\mathrm{I}} \times \Theta^-_{\mathrm{I}} \tag{4.59}$$
$$\Theta_{\mathrm{III}} \to \Theta^{t=0}_{\mathrm{II}} \times_{\Theta^+_{\mathrm{I}} \times \Theta^-_{\mathrm{I}}} \Theta^{t=1}_{\mathrm{II}} \tag{4.60}$$
$$\Theta_{\mathrm{IV}} \to \Theta^{02}_{\mathrm{II}} \times_{\Theta^0_{\mathrm{I}} \times \Theta^2_{\mathrm{I}}} (\Theta^{01}_{\mathrm{II}} \times_{\Theta^1_{\mathrm{I}}} \Theta^{12}_{\mathrm{II}}) \tag{4.61}$$

This completes the definition of Θ_*.

Definition 4.35 (Virtual moduli counts)**.** An element $\theta \in \Theta_*$ evidently gives rise to a morphism of $\mathcal{S}_*$-modules $p_* \circ \tilde{w}_* : \mathbb{Q}[\mathcal{S}_*] \to \mathbb{Q}$. Such a morphism corresponds to virtual moduli counts satisfying the relevant master equation by Remark 4.19.

Lemma 4.36. *Suppose $\overline{\mathcal{M}}_*(T)$ is regular and* $\mathrm{vdim}(T) = 0$. *Then $\overline{\mathcal{M}}_*(T) = \mathcal{M}_*(T)$ and $\#\overline{\mathcal{M}}_*(T)^{\mathrm{vir}}_\theta = \#\mathcal{M}_*(T)$.*

Proof. There are no nontrivial (effective) $T' \to T$ for dimension reasons, so $\overline{\mathcal{M}}_*(T) = \mathcal{M}_*(T)$. Now let us evaluate the diagram (4.58) at T and take cohomology to get:

$$\begin{tikzcd} \mathfrak{o}_T \arrow[r, "\tilde{w}_*"] \arrow[d, "w_*"] & \mathbb{Q} \arrow[r, "\mathrm{id}"] \arrow[d, "\mathrm{id}"] & \mathbb{Q} \\ \check{H}^\bullet(\overline{\mathcal{M}}_*(T); \mathfrak{o}_{\overline{\mathcal{M}}_*(T)}) \arrow[r, "{[\overline{\mathcal{M}}_*(T)]^{\mathrm{vir}}}"] & \mathbb{Q} & \end{tikzcd} \tag{4.62}$$

Unravelling the isomorphism in Lemma 4.26, we see the left vertical map is just the tautological map to $\check{H}^0$ via the identification $\mathfrak{o}_{\overline{\mathcal{M}}_*(T)} = \mathfrak{o}_T$. The bottom horizontal map is by definition the virtual fundamental class $[\overline{\mathcal{M}}_*(T)]^{\mathrm{vir}}$ from [Par15, Definition 5.1.1]. Thus it suffices to show that the virtual fundamental class $[\overline{\mathcal{M}}_*(T)]^{\mathrm{vir}}$ coincides with the usual fundamental class given that $\overline{\mathcal{M}}_*(T) = \overline{\mathcal{M}}_*(T)^{\mathrm{reg}}$. For $* \in \{\mathrm{I}, \mathrm{II}\}$, this is just [Par15, Lemma 5.2.6].

[15]Note that the categories $\mathcal{S}_*$ are essentially small, so the collection of such diagrams forms a set.

For $* \in \{\mathrm{III}, \mathrm{IV}\}$, we cannot immediately cite [Par15, Lemma 5.2.6] since the notion of regularity in Theorems III, IV is less restrictive than the one in the rest of the paper (see Remark 2.23). It suffices, though, to observe that the implicit atlas on $\overline{\mathcal{M}}_*(T)$ remains an implicit atlas if we use this less restrictive notion of regularity. The only thing to check is that the manifold/openness/submersion axioms still hold, but this is trivial since all thickened moduli spaces consist only of the top stratum $\mathcal{M}_*(T)_I$, and hence can be described by the usual Banach manifold Fredholm setup. $\square$

Lemma 4.37. *The set Θ_{I} is non-empty (resp. (4.59)–(4.61) are surjective).*

Proof. Concretely, we must construct B_*, p_*, w_*, $\tilde{w}_*$ as in Definition 4.34. We will in fact give a construction of such data by *induction* on $T \in \mathcal{S}_*$ with respect to the partial order from Definition 4.30. Note that the inductive nature of our proof (in addition to being the only reasonable approach) implies surjectivity of (4.59)–(4.61) (as opposed to mere non-emptiness of Θ_{II}, Θ_{III}, Θ_{IV}).

Finite subatlases $B_*(T) \subseteq A_*(T)$ exist since each $\overline{\mathcal{M}}_*(T)$ is compact, and can be made $\mathrm{Aut}(T)$-invariant by taking a union of translates.

The existence of p_* and w_* follow from Lemmas 4.33 and 4.26 respectively. Note that both are proved by induction on T, so their use here is permissible.

To show the existence of $\tilde{w}_*$, we argue by induction on T. Note that $\mathbb{Q}[\mathcal{S}_*]$ is cofibrant by Lemma 4.29. For non-maximal T, cofibrancy of $\mathbb{Q}[\mathcal{S}_*]$ determines $\tilde{w}_*(T)$ uniquely and implies that the diagram still commutes. For maximal T, we are faced with the following lifting problem:

$$\begin{tikzcd}
\operatorname*{colim}_{\operatorname{codim}(T'/T)\geq 1} \mathbb{Q}[\mathcal{S}_*](T') \arrow[r, "\tilde{w}_*"] \arrow[d, tail] & C_\bullet^{\mathrm{cof}}(E_*)(T) \arrow[d, two heads, "\sim"', "(4.52)"] \\
\mathbb{Q}[\mathcal{S}_*](T) \arrow[r, "(4.39)\circ w_*"'] \arrow[ur, dashed, "\tilde{w}_*"] & C_\bullet(E_*)(T)
\end{tikzcd} \tag{4.63}$$

The left vertical map is injective by Lemma 4.29, and the right vertical map is a surjective quasi-isomorphism by Definition 4.31. All four complexes are bounded below. Now Lemma 4.38 ensures that a lift exists, and we make it $\mathrm{Aut}(T)$-equivariant by averaging. $\square$

Lemma 4.38. *Consider a diagram of chain complexes bounded below over a ring R:*

$$\begin{tikzcd}
A_\bullet \arrow[r] \arrow[d, tail] & X_\bullet \arrow[d, two heads, "\sim"] \\
B_\bullet \arrow[r] \arrow[ur, dashed] & Y_\bullet
\end{tikzcd} \tag{4.64}$$

where the right vertical map is a surjective quasi-isomorphism and the left vertical map is an injection whose cokernel is componentwise projective. Then there exists a lift as illustrated.

Proof. This is the fact that "cofibrations have the left lifting property with respect to acyclic fibrations" in the projective model structure on $\mathsf{Ch}_{\bullet\geq 0}(R)$. It can be proved by a straightforward diagram chase. $\square$

4.9 Symmetric monoidal structure on Θ_{I}, Θ_{II}

Proposition 4.39. *There are functorial maps:*

$$\Theta_{\mathrm{I}}(Y,\lambda,J) \leftarrow \prod_{i\in I}\Theta_{\mathrm{I}}(Y_i,\lambda_i,J_i) \qquad \textit{for } (Y,\lambda,J)=\bigsqcup_{i\in I}(Y_i,\lambda_i,J_i) \tag{4.65}$$

$$\Theta_{\mathrm{II}}(\hat{X},\hat{\omega},\hat{J}) \leftarrow \prod_{i\in I}\Theta_{\mathrm{II}}(\hat{X}_i,\hat{\omega}_i,\hat{J}_i) \qquad \textit{for } (\hat{X},\hat{\omega},\hat{J})=\bigsqcup_{i\in I}(\hat{X}_i,\hat{\omega}_i,\hat{J}_i) \tag{4.66}$$

preserving the virtual moduli counts.

Proof. Note that (restricting to effective objects):

$$\mathcal{S}_{\mathrm{I}}(Y,\lambda,J)=\bigsqcup_{i\in I}\mathcal{S}_{\mathrm{I}}(Y_i,\lambda_i,J_i) \qquad \text{for } (Y,\lambda,J)=\bigsqcup_{i\in I}(Y_i,\lambda_i,J_i) \tag{4.67}$$

$$\mathcal{S}_{\mathrm{II}}(\hat{X},\hat{\omega},\hat{J})=\bigsqcup_{i\in I}\mathcal{S}_{\mathrm{II}}(\hat{X}_i,\hat{\omega}_i,\hat{J}_i) \qquad \text{for } (\hat{X},\hat{\omega},\hat{J})=\bigsqcup_{i\in I}(\hat{X}_i,\hat{\omega}_i,\hat{J}_i) \tag{4.68}$$

so an $\mathcal{S}_{\mathrm{I}}(Y,\lambda,J)$-module is the same as a tuple of $\mathcal{S}_{\mathrm{I}}(Y_i,\lambda_i,J_i)$-modules for $i\in I$ (and the same for $\mathcal{S}_{\mathrm{II}}$-modules). There are also natural inclusions:

$$A_{\mathrm{I}}(Y,\lambda,J)(T) \hookleftarrow A_{\mathrm{I}}(Y_i,\lambda_i,J_i)(T) \qquad \text{for } T\in\mathcal{S}_{\mathrm{I}}(Y_i,\lambda_i,J_i) \tag{4.69}$$

$$A_{\mathrm{II}}(\hat{X},\hat{\omega},\hat{J})(T) \hookleftarrow A_{\mathrm{II}}(\hat{X}_i,\hat{\omega}_i,\hat{J}_i)(T) \qquad \text{for } T\in\mathcal{S}_{\mathrm{II}}(\hat{X}_i,\hat{\omega}_i,\hat{J}_i) \tag{4.70}$$

Precisely, these maps preserve r, E, D, and extend λ as zero on $Y\setminus Y_i$ (resp. $\hat{X}\setminus\hat{X}_i$).

Now the map (4.65) is defined by taking the images of the sets $B_{\mathrm{I}}(T)$ under (4.69) and using the "same" diagrams of $\mathcal{S}_{\mathrm{I}}$-modules. The map (4.66) is defined similarly. It follows by definition that these maps are functorial and preserve the virtual moduli counts. □

5 Gluing

This section is devoted to the proof of Theorems 3.23 and 3.24. Namely, we prove that the regular loci in the thickened moduli spaces $(\overline{\mathcal{M}}_{\mathrm{I}})_I$, $(\overline{\mathcal{M}}_{\mathrm{II}})_I$, $(\overline{\mathcal{M}}_{\mathrm{III}})_I$, $(\overline{\mathcal{M}}_{\mathrm{IV}})_I$ admit the expected local topological descriptions in terms of the spaces G_{I}, G_{II}, G_{III}, G_{IV} from §2.6, and we verify that the natural "geometric" and "analytic" maps between orientation lines agree. We will give the argument for all $*\in\{\mathrm{I},\mathrm{II},\mathrm{III},\mathrm{IV}\}$ simultaneously.

The result we prove here is very similar to the gluing theorem in Hutchings–Taubes [HT07, HT09a] (simpler, in fact, since we have no obstruction bundle); see also [Par15, §§B–C].

5.1 Gluing setup

Proof of Theorem 3.23. Fix $*\in\{\mathrm{I},\mathrm{II},\mathrm{III},\mathrm{IV}\}$, $T\in\mathcal{S}_*$, $I\subseteq J\subseteq\bar{A}_*(T)$, $T'\to T$, and $\tilde{x}_0\in\mathcal{M}_*(T')_J$ with $s_{J\setminus I}(\tilde{x}_0)=0$ and $\psi_{IJ}(\tilde{x}_0)\in\mathcal{M}_*(T')_I^{\mathrm{reg}}$. Denote by $x_0\in\overline{\mathcal{M}}_*(T)_J^{\mathrm{reg}}$ the image of $\tilde{x}_0$ under the inclusion $\overline{\mathcal{M}}_*(T')_J/\mathrm{Aut}(T'/T)\hookrightarrow\overline{\mathcal{M}}_*(T)_J$. Since $\mathrm{Aut}(T'/T)$ acts

freely on $\overline{\mathcal{M}}_*(T')_J$, the choice of lift $\tilde{x}_0$ of x_0 induces a stratification of a neighborhood of $x_0 \in \overline{\mathcal{M}}_*(T)_J$ by $(\mathcal{S}_*)_{T'//T}$.

Our goal is to construct a germ of homeomorphism:

$$\big((G_*)_{T'//T} \times E_{J\setminus I} \times \mathbb{R}^{\mu(T')-\#V_s(T')+\dim E_I}, (0,0,0)\big) \to \big(\overline{\mathcal{M}}_*(T)_J, x_0\big) \tag{5.1}$$

which lands in $\overline{\mathcal{M}}_*(T)_J^{\mathrm{reg}}$ and which commutes with the maps from both sides to $(\mathcal{S}_*)_{T'//T} \times \mathfrak{s}(T) \times E_{J\setminus I}$. We denote by $0 \in (G_*)_{T'//T}$ the basepoint corresponding to all gluing parameters equal to ∞ (i.e. corresponding to no gluing at all) and $t = t(x_0)$.

The basepoint $\tilde{x}_0$ corresponds to a map $u_0 : C_0 \to \hat{X}_0$ and an element $e_0 \in E_I \subseteq E_J$, along with "discrete data" consisting of asymptotic markers, matching isomorphisms, and markings $\phi_\alpha : (C_0)_\alpha \to \overline{\mathcal{M}}_{0,E^\pm(T_\alpha)\cup\{1,\ldots,r_\alpha\}}$ for $\alpha \in J$.

5.1.1 Domain stabilization via divisors $\hat{D}_{v,i}$ and points $q_{v,i}$

We first stabilize the domain C_0 by adding marked points $q_{v,i}$ where it intersects certain divisors $\hat{D}_{v,i}$, arguing much the same as we did in §3.5. Such divisors automatically stabilize the domains of all maps in a neighborhood of $x_0 \in \overline{\mathcal{M}}_*(T)_J$, and thus we need only consider stable domain curves in the main gluing argument which follows.

For every vertex $v \in V(T')$, we add marked points $q_{v,i} \in (C_0)_v$ and choose divisors (local codimension two submanifolds) $\hat{D}_{v,i} \subseteq (\hat{X}_0)_v$ (required to be $\mathbb{R}$-invariant if v is a symplectization vertex) with $u_0(q_{v,i}) \in \hat{D}_{v,i}$ intersecting transversally, such that $(C_0)_v$ equipped with the marked points $\{p_{v,e}\}_e$ and $\{q_{v,i}\}_i$ is stable.

To show the existence of such points, it suffices to show that each unstable irreducible component of $(C_0)_v$ (equipped with just $\{p_{v,e}\}_e$ as marked points) has a point (and hence a non-empty open set) where du_0 (resp. $\pi_\xi du_0$ if v is a symplectization vertex) is injective. If $(C_0)_v \not\subseteq (C_0)_\alpha$ for all $\alpha \in I$, then $u_0|(C_0)_v$ is $\hat{J}_v$-holomorphic and the existence of such points follows from the arguments given in the proofs of Lemmas 3.18 and 3.20. For components $(C_0)_v \subseteq (C_0)_\alpha$ for some $\alpha \in J$, such points exist since $u_v|(C_0)_v$ satisfies Definition 3.10(iii).

5.1.2 Glued cobordisms $\hat{X}_g$ and points q'_v

Given any[16] gluing parameter $g \in (G_*)_{T'//T}$, we may form the glued cobordism $\hat{X}_g$ as follows. Note that $\hat{X}_0$ is equipped with cylindrical coordinates (1.9)–(1.10) in each end. Now we truncate each positive (resp. negative) end $[0,\infty) \times Y_e$ (resp. $(-\infty, 0] \times Y_e$) to $[0, g_e] \times Y_e$ (resp. $[-g_e, 0] \times Y_e$) and identify truncated ends by translation by g_e (if $g_e = \infty$ we do nothing). By construction, $\hat{X}_g$ is the target for pseudo-holomorphic buildings of type $T' \to T'_g$ (i.e. the image of g under the map $(G_*)_{T'//T} \to (\mathcal{S}_*)_{T'//T}$).

For every symplectization vertex $v \in V_s(T')$, fix a section q'_v of the universal curve over a neighborhood of $C_0 \in \overline{\mathcal{M}}_{0,E^\pm(T')\cup\{q_{v,i}\}_{v,i}}$ such that $q'_v(C_0) \in (C_0)_v$. Now over a neighborhood of $x_0 \in \overline{\mathcal{M}}_*(T)_J$, the sections q'_v determine poinits q'_v in the domain of each map. Over this neighborhood, we can identify the target with $\hat{X}_g$ for a unique $g \in (G_*)_{T'//T}$ by requiring

[16]All statements involving a choice of $g \in (G_*)_{T'//T}$ carry the (often tacit) assumption that g lies in a sufficiently small neighborhood of 0.

that each q'_v be mapped to the corresponding "zero level" $0_v \subseteq \hat{X}_g$ (i.e. the descent of $0_v := \{0\} \times Y_v \subseteq \hat{X}_0$).

5.1.3 Family of almost complex structures j_y on C_0

We now proceed to fix a family of almost complex structures on C_0 inducing a diffeomorphism onto the relevant moduli space of marked Riemann surfaces. Let $N \subseteq C_0$ denote the set of nodes (excluding $\{p_{v,e}\}_{v,e}$), and let $N_v := N \cap (C_0)_v$.

For every vertex $v \in V(T')$, fix a linear map:

$$A : \mathbb{C}^{2\#\{p_{v,e}\}_e + \#\{q_{v,i}\}_i - \#N_v - 3} \to C_c^\infty((C_0)_v \setminus (\{q_{v,i}\}_i \cup \{p_{v,e}\}_e \cup N_v), \mathrm{End}^{0,1}(T(C_0)_v)) \tag{5.2}$$

inducing an isomorphism onto the tangent space to $\overline{\mathcal{M}}^{\#\text{nodes}=\#N_v}_{0,\#\{q_{v,i}\}_i+(\#\{p_{v,e}\}_e)_{(2)}}$ at $(C_0)_v$ equipped with its marked points $\{p_{v,e}\}_e$ and $\{q_{v,i}\}_i$. Denote by j_0 the almost complex structure on C_0. Let:

$$\mathcal{J}_v := \mathbb{C}^{2\#\{p_{v,e}\} + \#\{q_{v,i}\}_i - \#N_v - 3} \qquad \mathcal{J} := \prod_{v \in V(T')} \mathcal{J}_v \tag{5.3}$$

For any[17] $y \in \mathcal{J}_v$, let $j_y := j_0 e^{A(y)}$ (an almost complex structure on $(C_0)_v$). This induces a local diffeomorphism:

$$\mathcal{J}_v \to \overline{\mathcal{M}}^{\#\text{nodes}=\#N_v}_{0,\{q_{v,i}\}_i+(\{p_{v,e}\}_e)_{(2)}} \tag{5.4}$$

5.1.4 Cylindrical coordinates on C_0 and $\hat{X}_0$

We now fix positive (resp. negative) holomorphic cylindrical coordinates:

$$[0, \infty) \times S^1 \to C_0 \tag{5.5}$$
$$(-\infty, 0] \times S^1 \to C_0 \tag{5.6}$$

near each positive (resp. negative) puncture. We also fix such cylindrical coordinates on either side of each node $n \in N$ (choosing which side is positive/negative arbitrarily). We assume that with respect to these fixed cylindrical coordinates, we have:

$$u_0(s,t) = (Ls, \tilde{\gamma}(t)) + o(1) \quad \text{as } |s| \to \infty \tag{5.7}$$

(i.e. the constant b in (2.1) vanishes) with respect to the cylindrical coordinates (1.9)–(1.10) on $\hat{X}_0$. To ensure that we can achieve (5.7), we allow the possibility that (5.5)–(5.6) are only defined for $|s|$ sufficiently large.

5.1.5 Glued curves C_g and points q''_e

Given any[18] gluing parameter $\alpha \in \mathbb{C}^{E^{\text{int}}(T')} \times \mathbb{C}^N$ (i.e. one for each interior edge of T' and one for each node $n \in N$), we may form the glued curve C_α as follows. For each interior

[17] All statements involving a choice of $y \in \mathcal{J}_v$ carry the (often tacit) assumption that y lies in a sufficiently small neighborhood of 0.

[18] All statements involving a choice of $\alpha \in \mathbb{C}^{E^{\text{int}}(T')} \times \mathbb{C}^N$ carry the (often tacit) assumption that α lies in a sufficiently small neighborhood of 0.

edge $v \xrightarrow{e} v'$ (or node $n \in N$), we truncate the positive (resp. negative) end $[0, \infty) \times S^1$ (resp. $(-\infty, 0] \times S^1$) to $[0, S] \times S^1$ (resp. $[-S, 0] \times S^1$) and identify them by $s = s' + S$ and $t = t' + \theta$ where $\alpha = e^{-S+i\theta}$ (if $\alpha = 0$ we do nothing). Note that the points $q_{v,i}$ and the complex structures j_y both descend naturally to both C_α. This operation gives local diffeomorphisms:

$$\mathcal{J}_v \times \mathbb{C}^{N_v} \to \overline{\mathcal{M}}_{0,\{q_{v,i}\}_i+(\{p_{v,e}\}_e)_{(2)}} \tag{5.8}$$

$$\mathcal{J} \times \mathbb{C}^N \to \prod_{v \in V(T)} \overline{\mathcal{M}}_{0,\{q_{v,i}\}_i+(\{p_{v,e}\}_e)_{(2)}} \tag{5.9}$$

The glued curve C_α is actually a bit too general for our purposes: in the current setting the gluing parameters at interior edges must come from the same gluing parameters $g \in (G_*)_{T'//T}$ used to glue the target cobordism.

So, given a gluing parameter $g \in (G_*)_{T'//T} \times \mathbb{C}^N$, we define C_g as above, where the gluing parameters at interior edges are given by $S_e := L_e^{-1} g_e$ and the unique θ_e corresponding to the given matching isomorphism $S_{p_{v,e}} C_v \to S_{p_{v',e}} C_{v'}$. For e with $g_e < \infty$, let $q''_e \in C_g$ denote the point in the middle of the corresponding neck, namely $s = \pm\frac{1}{2}S$ with respect to the coordinates (5.5)–(5.6) (the angular coordinate $t \in S^1$ of q''_e is irrelevant as long as it is fixed). Now for any fixed $g \in (G_*)_{T'//T}$, this construction gives a map:

$$\mathcal{J} \times \mathbb{C}^N \to \prod_{v \in V(T'_g)} \overline{\mathcal{M}}_{0,\{q_{v',i}\}_{v' \mapsto v,i}+\{q''_e\}_{e \mapsto v}+(\{p_{v,e}\}_e)_{(2)}} \tag{5.10}$$

Moreover, this is a diffeomorphism (onto its image) when restricted to some neighborhood of $0 \in \mathcal{J} \times \mathbb{C}^N$ of uniform size independent of g near zero.

The purpose of including the points q''_e as part of the glued data is precisely so that (5.10) is a local diffeomorphism. We will impose a codimension two condition on maps at q''_e, so that over a neighborhood of $x_0 \in \overline{\mathcal{M}}_*(T)_J$, the point q''_e in the domain is determined uniquely, and hence the parameters $(y, \alpha) \in \mathcal{J} \times \mathbb{C}^N$ are determined uniquely.

We denote by $q'_v \in (C_g, j_y)$ the value of the section q'_v at (C_g, j_y) for $g \in (G_*)_{T'//T} \times \mathbb{C}^N$ and $y \in \mathcal{J}$ (note that this q'_v may *not* coincide with the descent of $q'_v \in C_0$, even for $g = 0$).

5.1.6 Preglued maps u_g

We now define a "preglued" map $u_g : C_g \to \hat{X}_g$ close to $u_0 : C_0 \to \hat{X}_0$. As we shall see later, this preglued map is very close to solving the relevant (thickened) pseudo-holomorphic curve equation. Our goal will then be to understand the true solutions near u_g and to show that this construction gives a local parameterization of the moduli space near u_0.

Fix a smooth (cutoff) function $\chi : \mathbb{R} \to [0, 1]$ satisfying:

$$\chi(x) = \begin{cases} 1 & x \leq 0 \\ 0 & x \geq 1 \end{cases} \tag{5.11}$$

Definition 5.1 (Flattening)**.** For $g \in (G_*)_{T'//T} \times \mathbb{C}^N$, we define the "flattened" map:

$$u_{0|g} : C_0 \to \hat{X}_0 \tag{5.12}$$

as follows. Away from the ends, $u_{0|g}$ coincides with u_0. Over a positive end asymptotic to a parameterized Reeb orbit $\tilde{\gamma}(t) := u_0(\infty, t)$, we define $u_{0|g}$ as:

$$u_{0|g}(s,t) := \begin{cases} u_0(s,t) & s \leq \frac{1}{6}S \\ \exp_{(Ls,\tilde{\gamma}(t))}\Big[\chi\big(s - \frac{1}{6}S\big) \cdot \exp^{-1}_{(Ls,\tilde{\gamma}(t))} u_0(s,t)\Big] & \frac{1}{6}S \leq s \leq \frac{1}{6}S + 1 \\ (s, \tilde{\gamma}(t)) & \frac{1}{6}S + 1 \leq s \end{cases} \tag{5.13}$$

Over a positive end at a node $n \in N$, we define $u_{0|g}$ as:

$$u_{0|g}(s,t) := \begin{cases} u_0(s,t) & s \leq \frac{1}{6}S \\ \exp_{u_0(n)}\Big[\chi\big(s - \frac{1}{6}S\big) \cdot \exp^{-1}_{u_0(n)} u_0(s,t)\Big] & \frac{1}{6}S \leq s \leq \frac{1}{6}S + 1 \\ (s, \tilde{\gamma}(t)) & \frac{1}{6}S + 1 \leq s \end{cases} \tag{5.14}$$

An analogous definition applies over negative ends. Here $\exp : T\hat{X}_0 \to \hat{X}_0$ denotes any fixed exponential map (i.e. a smooth map defined in a neighborhood of the zero section satisfying $\exp(p,0) = p$ and $d\exp(p,\cdot) = \mathrm{id}_{T_p\hat{X}_0}$) which is $\mathbb{R}$-equivariant in any end.

Definition 5.2 (Pregluing)**.** For $g \in (G_*)_{T'//T} \times \mathbb{C}^N$, we define the "preglued" map:

$$u_g : C_g \to \hat{X}_g \tag{5.15}$$

as the obvious "descent" of $u_{0|g}$ from C_0 to C_g.

5.2 Gluing estimates

With the above setup understood, our aim is now to describe the "true solutions" close to the "approximate solution" $u_g : C_g \to \hat{X}_g$. This forms the core part of the gluing argument.

5.2.1 Weighted Sobolev norms

Our first step is simply to fix norms on the Sobolev spaces $W^{k,2,\delta}$ relevant for us. More precisely, what is important is the choice of norms *up to commensurability uniform in g near zero*, since all of the key gluing estimates must be uniform in the limit $g \to 0$. Fix metrics and connections as in Definition 2.17 on C_0 and $\hat{X}_0$ with respect to the ends fixed above and which (for convenience) agree across the parts to be glued (thus descending to C_g and $\hat{X}_g$).

Away from the ends/necks of C_g, we use the usual $W^{k,2}$-norm. In an end asymptotic to a Reeb orbit, the contribution to the norm squared is:

$$\int_{[0,\infty)\times S^1} \sum_{j=0}^{k} \big|D^j f\big|^2 e^{2\delta s}\, ds\, dt \tag{5.16}$$

In an end asymptotic to a node $n \in N$, we distinguish two cases: for Sobolev spaces of $(0,1)$-forms on C_g, we use (5.16), and for spaces of functions we allow decay to a constant, i.e. the contribution to the norm squared is:

$$|f(n)|^2 + \int_{[0,\infty)\times S^1} \sum_{j=0}^{k} \big|D^j[f - f(n)]\big|^2 e^{2\delta s}\, ds\, dt \tag{5.17}$$

The contribution of a neck is given as follows (first is for necks over Reeb orbits and necks over nodes for spaces of $(0,1)$-forms; second is for necks over nodes for spaces of functions):

$$\int_{[0,S]\times S^1} \sum_{j=0}^{k} \left|D^j f\right|^2 e^{2\delta \min(s,S-s)}\, ds\, dt \tag{5.18}$$

$$\left|\frac{1}{2\pi}\int_{S^1} f\left(\tfrac{1}{2}S,t\right) dt\right|^2 + \int_{[0,S]\times S^1} \sum_{j=0}^{k} \left|D^j \left[f - \frac{1}{2\pi}\int_{S^1} f\left(\tfrac{1}{2}S,t\right) dt\right]\right|^2 e^{2\delta \min(s,S-s)}\, ds\, dt \tag{5.19}$$

We use parallel transport to make sense of the differences $f-f(n)$, etc. To meausre/differentiate $(0,1)$-forms, we use the usual Riemannian metric on $[0,\infty)\times S^1$ (resp. $[0,S]\times S^1$) and usual flat connection on forms.

Different choices of metrics and connections yield norms which are uniformly equivalent for any fixed $k \geq 0$ and admissible δ.

Remark 5.3. For the purposes of the gluing argument which follows, it is sufficient to work with some fixed choice of sufficiently large k and admissible $\delta > 0$. Nevertheless, we will try to be precise about exactly where these constraints on k and δ are needed. Also note that, although the choice of (k,δ) affects the constants appearing in most estimates, it does *not* affect any of the actual maps we will study and/or construct.

Note that we now put weights over the nodal ends, whereas earlier we defined the linearized operator (3.10) on Sobolev spaces without weights near the nodes. It is therefore crucial to observe that adding weights does not change the kernel or cokernel of the linearized operator. To verify that this indeed is the case, argue as follows (assuming $k \geq 2$ and admissible $\delta > 0$). Let us temporarily use $W^{k,2,\delta,\delta}$ to denote the Sobolev spaces defined above (namely with weights in both the Reeb and nodal ends) and $W^{k,2,\delta}$ to denote the Sobolev spaces from Definition 2.18 with weights only in the Reeb ends; thus there are dense inclusions $W^{k,2,\delta} \hookrightarrow W^{k,2,\delta,\delta}$. The map on kernels is obviously injective, and the map on cokernels is obviously surjective. To show surjectivity on kernels and injectivity on cokernels, it suffices to show that if $D\xi = \eta$ with $\xi \in W^{k,2,\delta,\delta}$ and $\eta \in W^{k-1,2,\delta}$, then $\xi \in W^{k,2,\delta}$. Now as distributions, we have $D\xi = \eta + \epsilon$, for some distribution ϵ supported over the inverse images of the nodes $\tilde{N}_v \subseteq \tilde{C}_v$. By elliptic regularity, it suffices to show that $\epsilon = 0$. But now for any smooth test function φ supported over a small neighborhood of $\tilde{N}_v$, we have:

$$\langle \epsilon, \varphi\rangle = \langle D\xi - \eta, \varphi\rangle = \langle \xi, D^*\varphi\rangle - \langle \eta, \varphi\rangle \tag{5.20}$$

The first term is bounded by $c \cdot \|\varphi\|_{1,1}$ since D^* (the formal adjoint) is first order and $\xi \in W^{k,2,\delta,\delta} \subseteq C^0$. The second term is bounded by $c \cdot \|\varphi\|_2$ since $\eta \in W^{k-1,2,\delta} \subseteq L^2$. Now ϵ is supported at a finite number of points, so must be a linear combination of δ-functions and their derivatives, but these do not satisfy such bounds (recall that $W^{1,1} \not\hookrightarrow C^0$ since we are in two dimensions).

5.2.2 Nonlinear Fredholm setup for fixed g

We now formulate precisely what we mean by "solutions close to $u_g : C_g \to \hat{X}_g$". What we mean is "small zeroes of $\mathcal{F}_g$", where $\mathcal{F}_g$ is the map:

$$\mathcal{F}_g : W^{k,2,\delta}(C_g, u_g^* T\hat{X}_g)_{\substack{\xi(q_{v,i}) \in T\hat{D}_{v,i} \\ \pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda}\xi(q_e'')=0}} \oplus \mathcal{J} \oplus E_J \\ \to W^{k-1,2,\delta}(\tilde{C}_g, u_g^*(T\hat{X}_g)_{\hat{J}_{g_t}} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}_g}) \oplus \mathbb{R}^{V_s(T')} \tag{5.21}$$

defined as follows:

$$\mathcal{F}_g(\xi, y, e) := \\ \Big[(\mathsf{PT}^{g_t}_{\exp_{u_g}\xi \to u_g} \otimes I_y) \Big(d(\exp_{u_g}\xi) + \sum_{\alpha \in J} \lambda_\alpha((e_0 + e)_\alpha)(\phi^\xi_\alpha, \exp_{u_g}\xi) \Big)^{0,1}_{j_y, \hat{J}_{g_t}} \\ \oplus \bigoplus_{v \in V_s(T')} \pi_{\mathbb{R}}(\exp_{u_g}\xi)(q_v'(y)) \Big] \tag{5.22}$$

We explain the notation. We denote by $\exp : T\hat{X}_0 \to \hat{X}_0$ a fixed exponential map which is $\mathbb{R}$-equivariant in ends and over symplectizations and agrees across the parts to be glued, thus descending to $\hat{X}_g$. We fix a $\hat{J}_0$-linear connection on $T\hat{X}_0$ which is $\mathbb{R}$-equivariant in ends and over symplectizations and agrees across the parts to be glued, thus descending to $\hat{X}_g$. We denote by PT^{g_t} parallel transport with respect to the $\hat{J}_{g_t}$-linear part of this fixed connection. The map $I_y : \Omega^{0,1}_{C_g,j_y} \to \Omega^{0,1}_{C_g,j_0}$ denotes the composition $\Omega^{0,1}_{C_g,j_y} \to \Omega^1_{C_g} \otimes_{\mathbb{R}} \mathbb{C} \to \Omega^{0,1}_{C_g,j_0}$ (which is $\mathbb{C}$-linear).

The map $\mathcal{F}_g$ is defined over the ball of some fixed radius $c'_{k,\delta} > 0$ uniformly in g near zero, for any $k \geq 3$ and admissible $\delta \geq 0$. The constraint $k \geq 3$ ensures that $W^{k,2} \hookrightarrow C^1$, which is needed so that ϕ^ξ_α defined for $\|\xi\|_{k,2,\delta}$ small.

5.2.3 Estimate for $\|\mathcal{F}_g(0)\|$

We now show that $\mathcal{F}_g(0)$ is very small (i.e. the preglued map $u_g : C_g \to \hat{X}_g$ is very close to being a true solution).

Lemma 5.4. *We have:*

$$\|\mathcal{F}_g(0)\|_{k-1,2,\delta} \to 0 \quad \textit{as } g \to 0 \tag{5.23}$$

for all $k \geq 1$ and admissible δ.

Proof. Away from the necks and ends, the 1-form part of $\mathcal{F}_g(0)$ is only nonzero because of using $\hat{J}_{g_t}$ in place of $\hat{J}_0$. This difference clearly goes to zero as $g \to 0$.

Over Reeb ends, the 1-form part of $\mathcal{F}_g(0)$ is identically zero. Over Reeb necks, the 1-form part of $\mathcal{F}_g(0)$ is supported near $\frac{1}{6}S$ and $\frac{5}{6}S$, and the desired estimate follows from the exponential convergence of (5.7) and the fact that δ is admissible.

The same applies to nodal ends/necks, except that in addition there is a term coming from using $\hat{J}_{g_t}$ in place of $\hat{J}_0$. This again is bounded as desired since $\delta < 1$.

The $\mathbb{R}^{V_s(T')}$ part of $\mathcal{F}_g(0)$ satisfies the desired estimate since $q_v' \in (C_g, j_0)$ approaches the descent of $q_v' \in (C_0, j_0)$ as $g \to 0$. $\square$

5.2.4 Regularity of the map $\mathcal{F}_g$

We now estimate the regularity of $\mathcal{F}_g$, i.e. we give uniform upper bounds on its derivatives near zero. This estimate is used when we apply the (Banach space) inverse function theorem to understand $\mathcal{F}_g^{-1}(0)$ near zero.

The first term in $\mathcal{F}_g$ (the usual pseudo-holomorphic curve equation) is smooth and local. The second term (the "thickening" terms λ_α) is non-local; its only non-smoothness comes from the association $\xi \mapsto \phi_\alpha^\xi$. It thus is C^ℓ as long as the function which assigns to ξ the set $(\exp_{u_g} \xi)^{-1}(\hat{D}_\alpha)$ is C^ℓ. By the inverse function theorem, this is the case whenever $W^{k,2} \hookrightarrow C^\ell$, which in turn holds whenever $k \geq \ell + 2$. The third term is also C^ℓ whenever $W^{k,2} \hookrightarrow C^\ell$ (these both come down to the fact that the evaluation map $W^{k,2}(C,X) \times C \to X$ is of class C^ℓ whenever $W^{k,2} \hookrightarrow C^\ell$).

The following "quadratic estimate" is the specific type of bound on the derivatives of $\mathcal{F}_g$ which we will use later. Of course, much more should be true, namely that $\mathcal{F}_g$ is uniformly C^ℓ for $k \geq \ell + 2$, but for simplicity we will only state what we need.

Proposition 5.5. *For $\|\zeta\|_{k,2,\delta}, \|\xi\|_{k,2,\delta} \leq c'_{k,\delta}$, we have:*

$$\left\|\mathcal{F}_g'(0,\xi) - \mathcal{F}_g'(\zeta,\xi)\right\|_{k-1,2,\delta} \leq c_{k,\delta} \cdot \|\zeta\|_{k,2,\delta} \|\xi\|_{k,2,\delta} \tag{5.24}$$

for constants $c_{k,\delta} < \infty$ and $c'_{k,\delta} > 0$ uniformly in g near 0, for all $k \geq 4$ and admissible $\delta \geq 0$.

Proof. This may be proved by a straightforward (albeit long) calculation, treating each of the three terms in (5.22) separately. It suffices to write down precisely the above argument that $\mathcal{F}_g$ is of class C^2, and observe that the resulting estimates are all uniform over g near 0. We omit the details. □

Note that integrating (5.24) from ξ_1 to ξ_2 gives:

$$\left\|D_g(\xi_1 - \xi_2) - (\mathcal{F}_g\xi_1 - \mathcal{F}_g\xi_2)\right\|_{k-1,2,\delta} \leq c_{k,\delta} \cdot \|\xi_1 - \xi_2\|_{k,2,\delta} \cdot \max(\|\xi_1\|_{k,2,\delta}, \|\xi_2\|_{k,2,\delta}) \tag{5.25}$$

for $\|\xi_1\|_{k,2,\delta}, \|\xi_2\|_{k,2,\delta} \leq c'_{k,\delta}$.

5.2.5 Bounded right inverses and kernel gluing I: relating D_0 and D_g

The final step in understanding $\mathcal{F}_g^{-1}(0)$ is to construct a sufficiently nice bounded right inverse Q_g for $D_g := \mathcal{F}_g'(0,\cdot)$. In particular, we will need show that $\|Q_g\|$ is bounded uniformly for g near 0, and that $\operatorname{im} Q_g$ "varies continuously" (in a sense which we will make precise) as g varies. We will also construct a natural "kernel gluing" isomorphism $\ker D_0 \to \ker D_g$.

To study the linearized operator D_g, and in particular to construct Q_g, we consider the

following diagram, which allows us to relate D_g to D_0.

$$\begin{tikzcd}
W^{k,2,\delta}(C_g, u_g^* T\hat{X}_g)_{\substack{\xi(q_{v,i})\in T\hat{D}_{v,i}\\ \pi_{\mathbb{R}\partial_s\oplus\mathbb{R}R_\lambda}\xi(q''_e)=0}} \oplus \mathfrak{J} \oplus E_J \arrow[r, "D_g"] & W^{k-1,2,\delta}(\tilde{C}_g, u_g^*(T\hat{X}_g)_{\hat{J}_{g_t}} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}_g}) \oplus \mathbb{R}^{V_s(T')} \arrow[d, equal] \\
W^{k,2,\delta}(C_g, u_g^* T\hat{X}_g)_{\xi(q_{v,i})\in T\hat{D}_{v,i}} \oplus \mathfrak{J} \oplus E_J \arrow[r, "D_g"] \arrow[u, "\text{calib}"'] & W^{k-1,2,\delta}(\tilde{C}_g, u_g^*(T\hat{X}_g)_{\hat{J}_{g_t}} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}_g}) \oplus \mathbb{R}^{V_s(T')} \arrow[d, "\text{break}"] \\
W^{k,2,\delta}(C_0, u_{0|g}^* T\hat{X}_0)_{\xi(q_{v,i})\in T\hat{D}_{v,i}} \oplus \mathfrak{J} \oplus E_J \arrow[r, "D_{0|g}"] \arrow[u, "\text{glue}"'] \arrow[d, leftrightarrow, "\mathsf{PT}"'] & W^{k-1,2,\delta}(\tilde{C}_0, u_{0|g}^*(T\hat{X}_0)_{\hat{J}_{g_t}} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}_0}) \oplus \mathbb{R}^{V_s(T')} \arrow[d, leftrightarrow, "\mathsf{PT}\circ\mathrm{id}^{1,0}"] \\
W^{k,2,\delta}(C_0, u_0^* T\hat{X}_0)_{\xi(q_{v,i})\in T\hat{D}_{v,i}} \oplus \mathfrak{J} \oplus E_J \arrow[r, "D_0"] & W^{k-1,2,\delta}(\tilde{C}_0, u_0^*(T\hat{X}_0)_{\hat{J}_0} \otimes_{\mathbb{C}} \Omega^{0,1}_{\tilde{C}_0}) \oplus \mathbb{R}^{V_s(T')}
\end{tikzcd} \tag{5.26}$$

The horizontal maps D are all uniformly bounded ($D_{0|g}$ is defined as D_0 except with $u_{0|g}$ in place of u_0). The maps PT are parallel transport with respect to the fixed connection on $T\hat{X}_0$; they are uniformly bounded, as is $\mathrm{id}^{1,0} : (T\hat{X}_0)_{\hat{J}_0} \to (T\hat{X}_0)_{\hat{J}_{g_t}}$.

Let us define the break map from (5.26). Fix a smooth function $\bar{\chi} : \mathbb{R} \to [0,1]$ such that:

$$\bar{\chi}(x) = \begin{cases} 1 & x \le -1 \\ 0 & x \ge +1 \end{cases} \qquad \bar{\chi}(x) + \bar{\chi}(-x) = 1 \tag{5.27}$$

Now $\mathrm{break}(\eta)$ is simply η except over the ends of C_0, where we define it to be:

$$\mathrm{break}(\eta)(s,t) := \begin{cases} \eta(s,t) & s \le \frac{1}{2}S - 1 \\ \bar{\chi}(s - \frac{1}{2}S) \cdot \eta(s,t) & \frac{1}{2}S - 1 \le s \le \frac{1}{2}S + 1 \\ 0 & \frac{1}{2}S + 1 \le s \end{cases} \tag{5.28}$$

Thus the “trace” of $\mathrm{break}(\eta)$ from C_0 to C_g (adding along fibers) is precisely η. The norm of break is uniformly bounded.

Let us define the glue map from (5.26). The map glue acts only on the vector field component (it acts identically on the other components). Away from the necks, we set $\mathrm{glue}(\xi) := \xi$, and in any particular neck $[0,S] \times S^1 \subseteq C_g$, we define (respectively for necks near Reeb orbits and necks near nodes $n \in N$):

$$\mathrm{glue}(\xi)(s,t) := \begin{cases} \xi(s,t) & s \le \frac{1}{3}S - 1 \\ \chi(s - \frac{2}{3}S)\xi(s,t) + \chi(\frac{2}{3}S - s')\xi(s',t') & \frac{1}{3}S - 1 \le s \le \frac{2}{3}S + 1 \\ \xi(s',t') & \frac{2}{3}S + 1 \le s \end{cases} \tag{5.29}$$

$$\mathrm{glue}(\xi)(s,t) := \begin{cases} \xi(s,t) & s \le \frac{1}{3}S - 1 \\ \xi(n) + \chi(s - \frac{2}{3}S)[\xi(s,t) - \xi(n)] & \\ \qquad + \chi(\frac{2}{3}S - s')[\xi(s',t') - \xi(n)] & \frac{1}{3}S - 1 \le s \le \frac{2}{3}S + 1 \\ \xi(s',t') & \frac{2}{3}S + 1 \le s \end{cases} \tag{5.30}$$

(noting the corresponding ends $(s,t) \in [0,\infty) \times S^1 \subseteq C_0$ and $(s',t') \in (-\infty,0] \times S^1 \subseteq C_0$, glued via $s = s' + S$ and $t = t' + \theta$). The norm of glue is uniformly bounded.

Let us define the calib map from (5.26). For every edge $e \in E^{\mathrm{int}}(T')$ with $g_e < \infty$, we consider the vector field $C_g \to u_g^* T\hat{X}_g$ given in this neck by:

$$\chi(s - \tfrac{2}{3}S)\chi(\tfrac{2}{3}S - s') \cdot \partial_s u_g \tag{5.31}$$

We denote by $\mathbf{X}$ the $\mathbb{C}$-span of these vector fields. Now we have:

$$W^{k,2,\delta}(C_g, u_g^* T\hat{X}_g)_{\xi(q_{v,i}) \in T\hat{D}_{v,i}} = W^{k,2,\delta}(C_g, u_g^* T\hat{X}_g)_{\substack{\xi(q_{v,i}) \in T\hat{D}_{v,i} \\ \pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda}\xi(q_e'')=0}} \oplus \mathbf{X} \tag{5.32}$$

and the map calib is simply the associated projection onto the first factor. The norm of the projection onto the second factor is uniformly bounded since $q_e'' \in [0,S] \times S^1$ is given the largest weight in (5.18) and $\delta > 0$. Thus calib is also uniformly bounded for $\delta > 0$.

This completes the definition of the maps in (5.26), all of which are uniformly bounded.

5.2.6 Bounded right inverses and kernel gluing II: estimates

The diagram (5.26) does not commute, but is very close to commuting for g close to zero, as the following estimates make precise.

Lemma 5.6. *We have the following estimates:*

$$\|\mathsf{PT} \circ D_0 - D_{0|g} \circ \mathsf{PT}\| \to 0 \tag{5.33}$$

$$\|(D_g \circ \mathrm{glue})(\xi) - \eta\| = o(1) \cdot \|\xi\| \quad \textit{for } \mathrm{break}(\eta) = D_{0|g}\xi \tag{5.34}$$

$$\|D_g \circ \mathrm{calib} - D_g\| \to 0 \tag{5.35}$$

as $g \to 0$, for any fixed $k \geq 2$ and admissible $\delta > 0$.

Proof. To prove (5.33), argue as follows. The first difference between the two operators is over the $[\frac{1}{6}S, \infty) \times S^1$ subset of some ends. In this region, both are linear differential operators, which we may write in local coordinates (s,t) on C_0 and exponential coordinates on the target near the asymptotic orbit or point. The desired bound then follows from the exponential convergence of (5.7) (near $n \in N$, observe that smoothness of u_0 implies decay of all derivatives like $e^{-|s|}$ in cylindrical coordinates). The second difference between the two operators is $\hat{J}_0$ vs $\hat{J}_{g_t}$, and this is also bounded as desired, since $\hat{J}_{g_t} \to \hat{J}_0$ in C^∞.

To prove (5.34), argue as follows. The difference is only nonzero over the $([\frac{1}{3}S - 1, \frac{1}{3}S] \cup [\frac{2}{3}S, \frac{2}{3}S + 1]) \times S^1$ subsets of each neck. By symmetry, we discuss only the $[\frac{2}{3}S, \frac{2}{3}S + 1] \times S^1$ part, where it equals $D_{0|g}(\chi(s - \frac{2}{3}S)\xi(s,t))$. Now we note that $D_{0|g}(\chi(s - \frac{2}{3}S)\xi(s,t))$ has $W^{k-1,2,\delta}(\tilde{C}_0)$-norm bounded by $\|\xi\|_{k,2,\delta}$. But we are interested in the $W^{k-1,2,\delta}(\tilde{C}_g)$-norm, where the weight is smaller by a factor of $e^{-\frac{1}{3}\delta S}$, giving the desired estimate since $\delta > 0$.

To prove (5.35), argue as follows. It suffices to show that $\|D_g(X(\xi))\| = o(1)\|\xi\|$, where $X(\xi) \in \mathbf{X}$ denotes $\xi - \mathrm{calib}(\xi)$, i.e. the second projection in (5.32). Note that $\xi(q_e'')$ (which determines $X(\xi)$) is bounded in norm by a constant times $e^{-\frac{1}{2}\delta S}\|\xi\|_{k,2,\delta}$ (i.e. $\|\xi\|_{k,2,\delta}$ divided by the weight in the middle of the neck). Now $D_g(X(\xi))$ is only nonzero over the $([\frac{1}{3}S - 1, \frac{1}{3}S] \cup [\frac{2}{3}S, \frac{2}{3}S + 1]) \times S^1$ subsets of each neck, where the weight is $e^{\frac{1}{3}\delta S}$. Its norm is thus bounded $e^{(\frac{1}{3} - \frac{1}{2})\delta S}\|\xi\|_{k,2,\delta}$, giving the desired result since $\delta > 0$. ∎

5.2.7 Bounded right inverses and kernel gluing III: goal

Recall that by assumption, D_0 is surjective and the natural projection $\ker D_0 \to E_{J\setminus I}$ is surjective; indeed, this is what it means for $\psi_{IJ}(\tilde{x}_0)$ to lie in $\mathcal{M}_*(T')_I^{\mathrm{reg}}$. Let Q_0 denote any bounded right inverse for D_0, meaning $D_0 Q_0 = \mathbf{1}$. Then we have a direct sum decomposition:

$$W^{k,2,\delta}(C_0, u_0^* T\hat{X}_0)_{\xi(q_{v,i})\in T\hat{D}_{v,i}} \oplus \mathfrak{J} \oplus E_J = \ker D_0 \oplus \operatorname{im} Q_0 \tag{5.36}$$

In fact, choosing a bounded right inverse Q_0 is equivalent to choosing a closed complement $\operatorname{im} Q_0$ of $\ker D_0$. The classical Banach space implicit function theorem (taking as input Q_0 and the quadratic estimate (5.24)) then implies that the map from $\mathcal{F}_0^{-1}(0)$ to $\ker D_0$ by projection along $\operatorname{im} Q_0$ is a local diffeomorphism near zero.

Our goal is to generalize this setup to g in a neighborhood of zero (using (5.26) and Lemma 5.6). Namely, we will construct a right inverse Q_g for D_g (equivalently, we will choose a complement $\operatorname{im} Q_g$ for $\ker D_g$), so we that have a direct sum decomposition:

$$W^{k,2,\delta}(C_g, u_g^* T\hat{X}_g)_{\substack{\xi(q_{v,i})\in T\hat{D}_{v,i} \\ \pi_{\mathbb{R}\partial_s\oplus\mathbb{R}R_\lambda}\xi(q_e'')=0}} \oplus \mathfrak{J} \oplus E_J = \ker D_g \oplus \operatorname{im} Q_g \tag{5.37}$$

The same implicit function theorem argument applies as long as $\|Q_g\|$ is bounded uniformly for g near zero. Note also that uniform boundedness of Q_g implies that both projections in (5.37) are uniformly bounded (since they are given by $\mathbf{1} - Q_g D_g$ and $Q_g D_g$ respectively).

Now to ensure that the individual parameterizations of $\mathcal{F}_g^{-1}(0)$ by K_g near zero fit together continuously as g varies, we also need to show that the direct sum decomposition (5.37) is "continuous in g" in some sense. Let us now describe more precisely the sense we mean. For some points $w_i \in C_0 \setminus (\{p_{v,e}\}_{v,e} \cup N)$, consider the linear functional:

$$L_0 : W^{k,2,\delta}(C_0, u_0^* T\hat{X}_0)_{\xi(q_{v,i})\in T\hat{D}_{v,i}} \oplus \mathfrak{J} \oplus E_J \to \left(\bigoplus_i T_{u_0(w_i)}\hat{X}_0 \oplus \mathfrak{J} \oplus E_J \right) \Big/ B \tag{5.38}$$

for some subspace B projecting trivially onto $E_{J\setminus I}$. Fix B so that $L_0|_{\ker D_0}$ is an isomorphism; this is possible since $\ker D_0 \to E_{J\setminus I}$ is surjective. Since $L_0|_{\ker D_0}$ is an isomorphism, we have a direct sum decomposition:

$$W^{k,2,\delta}(C_0, u_0^* T\hat{X}_0)_{\xi(q_{v,i})\in T\hat{D}_{v,i}} \oplus \mathfrak{J} \oplus E_J = \ker D_0 \oplus \ker L_0 \tag{5.39}$$

Now denote by:

$$L_g : W^{k,2,\delta}(C_g, u_g^* T\hat{X}_g)_{\substack{\xi(q_{v,i})\in T\hat{D}_{v,i} \\ \pi_{\mathbb{R}\partial_s\oplus\mathbb{R}R_\lambda}\xi(q_e'')=0}} \oplus \mathfrak{J} \oplus E_J \to \left(\bigoplus_i T_{u_g(w_i)}\hat{X}_g \oplus \mathfrak{J} \oplus E_J \right) \Big/ B \tag{5.40}$$

the "same" linear functional, where $w_i \in C_g$ denote the descents of $w_i \in C_0$, so that there is a natural identification $T_{u_g(w_i)}\hat{X}_g = T_{u_0(w_i)}\hat{X}_0$. We will show that $L_g|_{\ker D_g}$ is still an isomorphism, and hence there is a direct sum decomposition:

$$W^{k,2,\delta}(C_g, u_g^* T\hat{X}_g)_{\substack{\xi(q_{v,i})\in T\hat{D}_{v,i} \\ \pi_{\mathbb{R}\partial_s\oplus\mathbb{R}R_\lambda}\xi(q_e'')=0}} \oplus \mathfrak{J} \oplus E_J = \ker D_g \oplus \ker L_g \tag{5.41}$$

We will construct Q_g with $\operatorname{im} Q_g = \ker L_g$, i.e. the direct sum decompositions (5.37) and (5.41) coincide. We will also define natural "kernel gluing" isomorphisms $\ker D_0 \xrightarrow{\sim} \ker D_g$ which agree with $L_g^{-1} \circ L_0$.

5.2.8 Bounded right inverses and kernel gluing IV: construction

We now construct the right inverses Q_g and the kernel gluing isomorphisms $\ker D_0 \xrightarrow{\sim} \ker D_g$ satisfying the desired properties discussed above.

We first recall the following general construction, which allows one to upgrade an "approximate right inverse" into a (true) right inverse.

Definition 5.7. Let $D : X \to Y$ be a bounded linear map between Banach spaces, and let $T : Y \to X$ be an *approximate right inverse*, meaning that $\|\mathbf{1} - DT\| < 1$. Then there is a (necessarily unique) associated right inverse $Q : Y \to X$ with the same image $\operatorname{im} Q = \operatorname{im} T$, namely $Q := T(DT)^{-1}$, where $DT : Y \to Y$ is invertible by the geometric series $\sum_{k=0}^{\infty}(\mathbf{1} - DT)^k$. Moreover, we have (trivially) that $\|Q\| \le \|T\|\,(1 - \|\mathbf{1} - DT\|)^{-1}$.

To define the right inverse Q_g, first define an approximate right inverse T_g of D_g as the following composition of maps in (5.26):

$$T_g := \mathrm{calib} \circ \mathrm{glue} \circ \mathsf{PT} \circ Q_0 \circ \mathsf{PT} \circ \mathrm{id}^{1,0} \circ \mathrm{break} \tag{5.42}$$

where Q_0 denotes the fixed right inverse of D_0 defined by the property that $\operatorname{im} Q_0 = \ker L_0$. A consequence of the estimates (5.33)–(5.35) (expressing the fact that (5.26) almost commutes) is that $\|\mathbf{1} - D_g T_g\| \to 0$ as $g \to 0$ (see [Par15, Lemma B.7.6]). Let Q_g denote the associated true right inverse, which is uniformly bounded for g near zero (since all the maps in (5.26) are uniformly bounded). Note that $L_g \circ \mathrm{calib} \circ \mathrm{glue} \circ \mathsf{PT} = L_0$ by inspection, so $\operatorname{im} Q_0 = \ker L_0$ implies that $\operatorname{im} Q_g = \operatorname{im} T_g \subseteq \ker L_g$.

We define the kernel gluing isomorphism $\ker D_0 \xrightarrow{\sim} \ker D_g$ as the composition:

$$(\mathbf{1} - Q_g D_g) \circ \mathrm{calib} \circ \mathrm{glue} \circ \mathsf{PT} : \ker D_0 \to \ker D_g \tag{5.43}$$

Note that $L_g \circ (1 - Q_g D_g) \circ \mathrm{calib} \circ \mathrm{glue} \circ \mathsf{PT} = L_0$ by inspection, and hence (5.43) is injective. Now we have by definition that $\operatorname{ind} D_0 = \mu(T') - \#V_s(T') - \#N + \dim E_J$ and $\operatorname{ind} D_g = \mu(T'_g) - \#V_s(T') - \#N + \dim E_J$, where $T' \to T'_g$ denotes the image of g under the map $(G_*)_{T'//T} \to (\mathcal{S}_*)_{T'//T}$. These indices coincide as remarked in Definition 2.25, so (5.43) is an isomorphism as both D_0 and D_g are surjective. Since (5.43) is an isomorphism, so is $L_g|_{\ker D_g}$, and it thus follows that the inclusion $\operatorname{im} Q_g \subseteq \ker L_g$ is in fact an equality $\operatorname{im} Q_g = \ker L_g$.

Remark 5.8. It is possible to prove that (5.43) is surjective directly at the cost of proving a few more estimates. This thus gives an *a priori* proof that $\mu(T) = \mu(T')$ for $T' \to T$.

Let us sketch the argument. Given $\ell \in \ker D_g$, some *a priori* estimates show that in any neck, ℓ decays rapidly to a constant vector field tangent to the trivial cylinder; moreover, this constant vanishes in Reeb necks since $\pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda} \ell(q''_e) = 0$. It follows that we can apply an "ungluing" operation to produce a κ of commensurable norm $\|\kappa\| \asymp \|\ell\|$ with $\|D_0\kappa\| = o(1) \cdot \|\kappa\|$ and $L_0\kappa = L_g\ell$. Now we have $\|\ell - (\mathrm{calib} \circ \mathrm{glue} \circ \mathsf{PT})(\kappa)\| = o(1) \cdot \|\ell\|$ by explicit calculation, and it follows that the image of $(\mathbf{1} - Q_0 D_0)\kappa \in \ker D_0$ under (5.43) is within distance $o(1) \cdot \|\ell\|$ of ℓ. Since this holds for all $\ell \in \ker D_g$, we conclude that (5.43) is surjective.

5.3 Gluing map

We now define the gluing map and show that it is a germ of homeomorphism. This is the "endgame" of the gluing argument, where we deduce the desired results from the technical work performed above.

5.3.1 Definition of the gluing map

We first recall (following our sketch in §5.2.7) how our work above implies that $\mathcal{F}_g^{-1}(0)$ is a manifold near zero and that projection along $\operatorname{im} Q_g$ provides a diffeomorphism between it and $\ker D_g$ near zero.

We have fixed a right inverse Q_g for $D_g = \mathcal{F}_g'(0,\cdot)$ with $\|Q_g\|$ bounded uniformly for g near zero. Now it follows from (5.24) that Q_g is an approximate right inverse to $\mathcal{F}_g'$ over the ball of some radius $c_{k,\delta}' > 0$ (uniform in g). Hence over this ball of radius $c_{k,\delta}' > 0$, the operator $\mathcal{F}_g'$ is surjective, i.e. $\mathcal{F}_g$ is transverse to zero. By the Banach space implicit function theorem, it thus follows that $\mathcal{F}_g^{-1}(0)$ is a C^ℓ-submanifold (for $k \geq \ell + 2$) which is transverse to $\operatorname{im} Q_g$.

Let us now show that map $\ker D_g \to \mathcal{F}_g^{-1}(0)$ given by projection along $\operatorname{im} Q_g$ (is well-defined and) is a diffeomorphism near zero. The key point is that the map $\mathbf{1} - Q_g\mathcal{F}_g$ is a contraction mapping when restricted to any slice $(\xi + \operatorname{im} Q_g) \cap B(c_{k,\delta}')$ with ξ sufficiently small in terms of $c_{k,\delta}'$. This follows from (5.25) and (5.23), which moreover imply that the contraction constant approaches zero (uniformly in g) as $c_{k,\delta}' \to 0$. This gives the desired result, and moreover shows that the projection along $\operatorname{im} Q_g$ to $\mathcal{F}_g^{-1}(0)$ is given (over the whole ball $B(c_{k,\delta}')$) by the limit of the Newton–Picard iteration $\xi \mapsto \xi - Q_g\mathcal{F}_g\xi$.

We can now define the gluing map, by precomposing the above local diffeomorphisms $\ker D_g \to \mathcal{F}_g^{-1}(0)$ with the kernel gluing isomorphisms $\ker D_0 \xrightarrow{\sim} \ker D_g$ and letting g vary. In other words, the gluing map:

$$\Big((G_*)_{T'//T} \times \mathbb{C}^N \times \ker D_0, (0,0)\Big) \to \Big(\overline{\mathcal{M}}_*(T)_J, x_0\Big) \tag{5.44}$$

sends (g,κ) to the map $\exp_{u_g} \kappa_g^\infty : C_g \to \hat{X}_g$, where $\kappa_g^\infty \in \mathcal{F}_g^{-1}(0)$ is unique intersection point $\mathcal{F}_g^{-1}(0) \cap (((\mathbf{1} - Q_gD_g) \circ \operatorname{calib} \circ \operatorname{glue} \circ \mathsf{PT})(\kappa) + \operatorname{im} Q_g)$ in $B(c_{k,\delta}')$. The discrete data for $\exp_{u_g} \xi : C_g \to \hat{X}_g$ is naturally inherited from that for $u_0 : C_0 \to \hat{X}_0$. Since $\mathcal{F}_g$ is transverse to zero at κ_g^∞, it follows that the image of the gluing map is contained in $\overline{\mathcal{M}}_*(T)_J^{\mathrm{reg}}$.

The gluing map evidently commutes with the maps from both sides to $(\mathcal{S}_*)_{T'//T} \times \mathfrak{s}(T) \times E_{J\setminus I}$ (recall that $\operatorname{im} Q_g = \ker L_g$ projects trivially onto $E_{J\setminus I}$ by definition).

Let us also note here that the inequality $\mu(T') - \#V_s(T') - 2\#N + \dim E_I \geq 0$ follows from the fact that D_0 is surjective, $\ker D_0 \to E_{J\setminus I}$ is surjective, and $\operatorname{ind} D_0 = \mu(T') - \#V_s(T') - 2\#N + \dim E_J$.

5.3.2 Properties of the gluing map

We now show that the gluing map is a germ of homeomorphism. More precisely, we show that the gluing map is continuous and that it is a germ of bijection. We then appeal to some point set topology to see that the gluing map is a germ of homeomorphism (though continuity of the inverse could also be proven directly).

Lemma 5.9. *The gluing map is continuous.*

Proof. Recall that κ_g^∞ may be described via the Newton–Picard iteration as follows. Namely, $\kappa_g^\infty = \lim_{i\to\infty} \kappa_g^i$, where:

$$\kappa_g^{i+1} = \kappa_g^i - Q_g \mathfrak{F}_g \kappa_g^i \tag{5.45}$$

$$\kappa_g^0 = (\mathrm{calib} \circ \mathrm{glue} \circ \mathsf{PT})(\kappa) \tag{5.46}$$

Note that there is no $\mathbf{1} - Q_g D_g$ in the definition of κ_g^0 (this is ok since $Q_g D_g \kappa_g^0 \in \operatorname{im} Q_g$).

Now suppose $(g_i, \kappa_i) \to (g, \kappa)$ (a convergent net), and let us show that $\exp_{u_{g_i}} (\kappa_i)_{g_i}^\infty : C_{g_i} \to \hat{X}_{g_i}$ approaches $\exp_{u_g} \kappa_g^\infty : C_g \to \hat{X}_g$ in the Gromov topology.

First, we claim that $\|(\kappa_i)_{g_i}^\infty - \kappa_{g_i}^\infty\|_{k,2,\delta} \to 0$. By uniform convergence of the Newton–Picard iteration, it suffices to show that $\|(\kappa_i)_{g_i}^n - \kappa_{g_i}^n\|_{k,2,\delta} \to 0$ for all n. The case $n = 0$ follows from uniform boundedness of $\mathrm{calib} \circ \mathrm{glue} \circ \mathsf{PT}$. The desired claim then follows by induction on n using (5.24). Now the claim implies that $\|(\kappa_i)_{g_i}^\infty - \kappa_{g_i}^\infty\|_\infty \to 0$, and thus it suffices to show that:

$$\exp_{u_{g_i}} \kappa_{g_i}^\infty : C_{g_i} \to \hat{X}_{g_i} \text{ approaches } \exp_{u_g} \kappa_g^\infty : C_g \to \hat{X}_g \tag{5.47}$$

Define $(\kappa_g^\infty)_{g_i}^0$ by (as the notation suggests) pregluing κ_g^∞ from C_g to C_{g_i} as follows. In any neck of C_{g_i} corresponding to a pair of ends of C_g, we preglue via $\mathrm{calib} \circ \mathrm{glue} \circ \mathsf{PT}$ as before (this operation is local to the ends/neck). In any neck of C_{g_i} corresponding to a neck of C_g, we simply use parallel transport and an nice diffeomorphism between the two necks (say, converging to the identity map in the C^∞ topology as $g_i \to g$). We may assume without loss of generality that there are no pairs of ends of C_{g_i} corresponding to a neck of C_g.

Now we claim that:

$$\|\mathfrak{F}_{g_i}((\kappa_g^\infty)_{g_i}^0)\|_{k-1,2,\delta} \to 0 \tag{5.48}$$

Away from the necks/ends, the 1-form part of $\mathfrak{F}_{g_i}((\kappa_g^\infty)_{g_i}^0)$ is nonzero only because of using $\hat{J}_{(g_i)_t}$ in place of $\hat{J}_{g_t}$. We have $\hat{J}_{(g_i)_t} \to \hat{J}_{g_t}$, so the desired estimate follows since $(\kappa_g^\infty)_{g_i}^0 = \kappa_g^\infty$ away from the ends/necks. Over the Reeb ends of C_{g_i}, the 1-form part vanishes. Over the Reeb necks of C_{g_i} corresponding to necks of C_g, the 1-form part approaches zero. Over the Reeb necks of C_{g_i} corresponding to pairs of ends of C_g, the estimate follows from the exponential decay of κ_g^∞ and u_0 and the fact that δ is admissible. The same applies to nodal necks, in addition considering the convergence $\hat{J}_{(g_i)_t} \to \hat{J}_{g_t}$. The $\mathbb{R}^{V_s(T')}$ part clearly approaches zero. This proves (5.48).

Now we consider the Newton–Picard iteration starting at $(\kappa_g^\infty)_{g_i}^0$, with limit $(\kappa_g^\infty)_{g_i}^\infty \in \mathfrak{F}_{g_i}^{-1}(0)$. By uniform contraction of the iteration and (5.48), we conclude that $\|(\kappa_g^\infty)_{g_i}^0 - (\kappa_g^\infty)_{g_i}^\infty\|_{k,2,\delta} \to 0$. It thus follows that:

$$\exp_{u_{g_i}} (\kappa_g^\infty)_{g_i}^\infty : C_{g_i} \to \hat{X}_{g_i} \text{ approaches } \exp_{u_g} \kappa_g^\infty : C_g \to \hat{X}_g \tag{5.49}$$

Now we claim that $(\kappa_g^\infty)_{g_i}^\infty = \kappa_{g_i}^\infty$, which is clearly sufficient to conclude the proof.

By construction, we have $L_{g_i}((\kappa_g^\infty)_{g_i}^\infty) = L_{g_i}((\kappa_g^\infty)_{g_i}^0) = L_g(\kappa_g^\infty) = L_g(\kappa_g^0) = L_0(\kappa)$ and similarly $L_{g_i}(\kappa_{g_i}^\infty) = L_{g_i}(\kappa_{g_i}^0) = L_0(\kappa)$. Thus $(\kappa_g^\infty)_{g_i}^\infty$ and $\kappa_{g_i}^\infty$ differ by an element of $\ker L_{g_i} = \operatorname{im} Q_{g_i}$, which is enough. ∎

Lemma 5.10. *The gluing map (5.44) is a germ of bijection. That is, for every sufficiently small neighborhood U of the basepoint in the domain, there exists an open neighborhood V of the basepoint in the target such that every $v \in V$ has a unique inverse image $u \in U$.*

Proof. Let $x \in \overline{\mathcal{M}}_*(T)_J$, and denote the corresponding map by $u : C \to \hat{X}$. We assume that x is sufficiently close to x_0, and we will show that x has a unique inverse image under the gluing map (5.44).

Concretely, x close to x_0 in the Gromov topology means the following. We may identify C with (C_β, j_w), for arbitrarily small gluing parameters $\beta \in \mathbb{C}^{E^{\mathrm{int}}(T')} \times \mathbb{C}^N$ and an almost complex structure j_w which agrees with j_0 except over a compact set (away from the ends/necks) where it is arbitrarily C^∞ close to j_0. Furthermore, the map $u : (C_\beta, j_w) \to \hat{X}$ is arbitrarily C^0-close to u_0 away from the Reeb ends/necks, and over the Reeb ends/necks it is arbitrarily C^0-close after forgetting the $\mathbb{R}$-coordinate. Note that by Lemma 5.11 and Arzelà–Ascoli, we in fact have u is arbitrarily C^∞-close to u_0 away from the ends/necks.

Now observe that there are unique points $q_{v,i} \in C$ close to $q_{v,i} \in C_0$ where u intersects $D_{v,i}$ transversally. These points give rise to unique points $q'_v \in C$ according to the sections chosen in §5.1.2. Now regarding $u(q'_v) \in \hat{X}$ as lying on the "zero section" determines (uniquely) a gluing parameter $g \in (G_*)_{T'//T}$ and an isomorphism $\hat{X} = \hat{X}_g$. Thus x corresponds to a map $u : (C_\beta, j_w) \to \hat{X}_g$.

Now in any Reeb end $[0, \infty) \times S^1 \subseteq C_\beta$ or Reeb neck $[0, S] \times S^1 \subseteq C_\beta$, we apply [HWZ02, Theorem 1.3] to see that u decays exponentially to a trivial cylinder. In particular, we conclude that for unglued edges, the tangent space marking of C_β at $\{p_{v,e}\}$ induced by the cylindrical coordinates on $\hat{X}_g$ is arbitrarily close to the tangent space marking descended from C_0. We also conclude that the gluing parameters $\beta \in \mathbb{C}^{E^{\mathrm{int}}(T')}$ are given by $L_e^{-1} g_e + o(1)$ and $\theta_e + o(1)$ (i.e. are very close to those coming from g), and we recover the point $q''_e \in C_\beta$ as the inverse image of $\exp_{u_{0|g}(q''_e)}(\ker \pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda})$. From these estimates, we conclude that x corresponds to a map $u : (C_{g,\alpha}, j_y) \to \hat{X}_g$ which is arbitrarily C^∞-close to $u_{g,\alpha}$ (including over the ends/necks, with respect to the cylindrical coordinates) respecting the tangent markings at $p_{v,e}$ and sending q''_e to $\exp_{u_{0|g}(q''_e)}(\ker \pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda})$. Here α and y can be assumed arbitrarily small, and we use the fact that (5.10) is a local diffeomorphism uniformly in g.

The injectivity radius of our fixed exponential map on $\hat{X}_g$ is bounded below uniformly, so x corresponds uniquely to some pair $(u = \exp_{u_{g,\alpha}} \xi : (C_{g,\alpha}, j_y) \to \hat{X}_g, e)$, where:

$$(\xi, y, e - e_0) \in W^{k,2,\delta}(C_{g,\alpha}, u^*_{g,\alpha} T\hat{X}_g)_{\substack{\xi(q_{v,i}) \in T\hat{D}_{v,i} \\ \pi_{\mathbb{R}\partial_s \oplus \mathbb{R}R_\lambda} \xi(q''_e) = 0}} \oplus \mathcal{J} \oplus E_J \tag{5.50}$$

has arbitrarily small norm due to the *a priori* estimates on u from Lemmas 5.11, 5.12, and [HWZ02, Theorem 1.3]. Now we observed in the gluing construction that for any fixed sufficiently small g, the gluing map gives a local diffeomorphism between $\ker D_0$ and $\mathcal{F}_g^{-1}(0)$, over a ball of size uniformly bounded below. Thus our map is uniquely in the image of the gluing map. ☐

Since the target of the gluing map is Hausdorff and the domain locally compact Hausdorff, it follows from continuity (Lemma 5.9) and bijectivity in a small neighborhood (Lemma 5.10) that the gluing map (5.44) is in fact a local homeomorphism, thus completing the proof of Theorem 3.23. ☐

5.4 Orientations

We now prove the compatibility of the geometric and analytic maps on orientation lines, namely that (3.14) commutes. We rely heavily on the gluing construction in §§5.1–5.3.

Proof of Theorem 3.24. The gluing map (5.44) (in the case $I = J$) allows us to describe the left vertical "geometric" map in (3.14) (i.e. the map induced by the topological structure of $\overline{\mathcal{M}}_*(T)_I^{\mathrm{reg}}$) near the basepoint x_0 as follows. Note that we may assume (for convenience) that $N = \varnothing$, since this locus is dense. Recall that there is a canonical identification:

$$\mathfrak{o}_{\ker D_0} = \mathfrak{o}^0_{T'} \otimes (\mathfrak{o}^\vee_{\mathbb{R}})^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I} \tag{5.51}$$

Now consider sufficiently small (g, κ), where $g \in (G_*)_{T'//T}$ lies in the top stratum $(T' \to T)$. Now $\mathcal{F}_g^{-1}(0)$ is a submanifold with tangent space $\ker \mathcal{F}'_g(\kappa_g^\infty, \cdot)$, and there is a canonical identification:

$$\mathfrak{o}_{\ker \mathcal{F}'_g(\kappa_g^\infty, \cdot)} = \mathfrak{o}^0_{T} \otimes (\mathfrak{o}^\vee_{\mathbb{R}})^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I} \tag{5.52}$$

The gluing map is differentiable with respect to κ since $\mathcal{F}_g^{-1}(0)$ is a submanifold transverse to $\operatorname{im} Q_g$. Its derivative is clearly given by the composition of $\ker D_0 \to \ker D_g$ and the map $\ker D_g \to \ker \mathcal{F}'_g(\kappa_g^\infty, \cdot)$ given by projecting off $\operatorname{im} Q_g$. This map $\ker D_0 \to \ker \mathcal{F}'_g(\kappa_g^\infty, \cdot)$ thus gives the "geometric" map $\mathfrak{o}^0_{T'} \to \mathfrak{o}^0_T$ when combined with the isomorphisms above.

Now the right inverse Q_g to $D_g = \mathcal{F}'_g(0, \cdot)$ is an approximate right inverse to $\mathcal{F}'_g(\xi, \cdot)$ for all $\xi \in B(c'_{k,\delta})$ by (5.24). Hence the kernel $\ker \mathcal{F}'_g(\xi, \cdot)$ forms a vector bundle over $B(c'_{k,\delta})$ which is canonically oriented by $\mathfrak{o}^0_T \otimes (\mathfrak{o}^\vee_{\mathbb{R}})^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I}$, and the map $\ker D_0 \to \ker \mathcal{F}'_g(\xi, \cdot)$ makes sense for all such ξ. Thus the geometric map on orientations is also given by the simpler map at $\xi = 0$:

$$\ker D_0 \xrightarrow{(1 - Q_g D_g) \circ \mathrm{cali} \circ \mathrm{glue} \circ \mathsf{PT}} \ker D_g \tag{5.53}$$

combined with the canonical identifications:

$$\mathfrak{o}_{\ker D_0} = \mathfrak{o}^0_{T'} \otimes (\mathfrak{o}^\vee_{\mathbb{R}})^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I} \tag{5.54}$$

$$\mathfrak{o}_{\ker D_g} = \mathfrak{o}^0_{T} \otimes (\mathfrak{o}^\vee_{\mathbb{R}})^{\otimes V_s(T')} \otimes \mathfrak{o}_{E_I} \tag{5.55}$$

Now this map (5.53) is precisely the sort of kernel pregluing map which defines the "analytic" map on orientations.

Strictly speaking, the analytic map $\mathfrak{o}^0_{T'} \to \mathfrak{o}^0_T$ is defined using a slightly different linearized operator (no $\mathcal{J}$, E_I, or point conditions), but this is only a "finite-dimensional" difference (note also that $\mathcal{J}$ is canonically oriented since it is a complex vector space). It is thus straightforward to relate them and see that they give rise to the same analytic map on orientations. □

5.5 Elliptic *a priori* estimates on pseudo-holomorphic curves

We record here some fundamental *a priori* estimates which guarantee the regularity of pseudo-holomorphic curves. These estimates play a fundamental role in the basic local properties of moduli spaces of pseudo-holomorphic curves. In particular, they are a crucial part of the proof of surjectivity of the gluing map in Lemma 5.10.

Lemma 5.11 (Gromov [Gro85])**.** *Let $u : D^2(1) \to (B^{2n}(1), J)$ be J-holomorphic, where J is tamed by $d\lambda$. For every $k < \infty$, we have:*

$$\|u\|_{C^k(D^2(1-s))} \le M \cdot (1 + s^{-k}) \tag{5.56}$$

for some constant $M = M(k, J, \lambda) < \infty$.

Proof. The Gromov–Schwarz Lemma (see Gromov [Gro85, 1.3.A] or Muller [Mul94, Corollary 4.1.4]) is the case $k = 1$. Standard elliptic bootstrapping allows one to upgrade this to bounds on all higher derivatives (see [Par15, Lemma B.11.4]). Note that it is enough to bound $D^k u$ at $0 \in D^2(1)$; the full result is recovered by restricting to the maximal disk centered at a given point in $D^2(1)$. □

Lemma 5.12 (Well-known)**.** *Let $u : [0, N] \times S^1 \to (B^{2n}(1), J)$ be J-holomorphic, where J is tamed by $d\lambda$. For every $k < \infty$ and $\epsilon > 0$, we have:*

$$\|u\|_{C^k([s,N-s]\times S^1)} \le M \cdot (1 + s^{-k}) e^{-(1-\epsilon)s} \tag{5.57}$$

for some constant $M = M(k, \epsilon, J, \lambda) < \infty$.

Proof. See [MS04] or [Par15, Proposition B.11.1]. □

A similar result for long J-holomorphic cylinders $u : [0, N] \times S^1 \to \mathbb{R} \times Y$ due to Hofer–Wysocki–Zehnder [HWZ02, Theorem 1.3] (see also [BEHWZ03, Proof of Proposition 5.7]) gives (under appropriate hypotheses) a bound of the form:

$$\|u(s,t) - (Ls + b, \tilde{\gamma}(t))\|_{C^k([s,N-s]\times S^1)} \le M \cdot e^{-(\delta_\gamma - \epsilon)\min(s,N-s)} \qquad s \ge s_0 \tag{5.58}$$

for some $b \in \mathbb{R}$, some $\tilde{\gamma} : S^1 \to Y$ with $\partial_t \tilde{\gamma} = L \cdot R_\lambda(\tilde{\gamma})$ parameterizing $\gamma \in \mathcal{P}$, and some constants $M, s_0 < \infty$ depending only on (Y, λ, J), the λ-energy of u, $k < \infty$, and $\epsilon > 0$.

References

[Abo14] Mohammed Abouzaid, *Symplectic cohomology and Viterbo's theorem*, Arxiv Preprint **arXiv:1312.3354v2** (2014), 1–184. 15

[AH09] Peter Albers and Helmut Hofer, *On the Weinstein conjecture in higher dimensions*, Comment. Math. Helv. **84** (2009), no. 2, 429–436. MR 2495800 (2010b:53152) 14

[AM12] Miguel Abreu and Leonardo Macarini, *Contact homology of good toric contact manifolds*, Compos. Math. **148** (2012), no. 1, 304–334. MR 2881318 15

[Aro57] N. Aronszajn, *A unique continuation theorem for solutions of elliptic partial differential equations or inequalities of second order*, J. Math. Pures Appl. (9) **36** (1957), 235–249. MR 0092067 (19,1056c)

[BC05] Frédéric Bourgeois and Vincent Colin, *Homologie de contact des variétés toroïdales*, Geom. Topol. **9** (2005), 299–313 (electronic). MR 2116317 (2005k:53166) 15

[BEE12] Frédéric Bourgeois, Tobias Ekholm, and Yasha Eliashberg, *Effect of Legendrian surgery*, Geom. Topol. **16** (2012), no. 1, 301–389, With an appendix by Sheel Ganatra and Maksim Maydanskiy. MR 2916289 12, 15

[BEHWZ03] F. Bourgeois, Y. Eliashberg, H. Hofer, K. Wysocki, and E. Zehnder, *Compactness results in symplectic field theory*, Geom. Topol. **7** (2003), 799–888. MR 2026549 (2004m:53152) 2, 3, 4, 5, 6, 16, 21, 22, 23, 74

[BEM14] Matthew Strom Borman, Yakov Eliashberg, and Emmy Murphy, *Existence and classification of overtwisted contact structures in all dimensions*, Arxiv Preprint **arXiv:1404.6157** (2014), 1–86. 13

[BH15] Erkao Bao and Ko Honda, *Definition of Cylindrical Contact Homology in dimension three*, Arxiv Preprint **arXiv:1412.0276** (2015), 1–53. 1

[BM04] Frédéric Bourgeois and Klaus Mohnke, *Coherent orientations in symplectic field theory*, Math. Z. **248** (2004), no. 1, 123–146. MR 2092725 (2005g:53173) 27, 29

[BN10] Frédéric Bourgeois and Klaus Niederkrüger, *Towards a good definition of algebraically overtwisted*, Expo. Math. **28** (2010), no. 1, 85–100. MR 2606237 (2011b:53212)

[BO09] Frédéric Bourgeois and Alexandru Oancea, *An exact sequence for contact- and symplectic homology*, Invent. Math. **175** (2009), no. 3, 611–680. MR 2471597 (2010e:53149) 15

[BO14] ———, *S^1-equivariant symplectic homology and linearized contact homology*, Arxiv Preprint **arXiv:1212.3731v3** (2014), 1–64. 15

[BO15] ———, *Erratum to: An exact sequence for contact- and symplectic homology*, Invent. Math. **200** (2015), no. 3, 1065–1076. MR 3348144 15

[Bou02] Frédéric Bourgeois, *A Morse–Bott approach to contact homology*, Ph.D. Thesis, Stanford University (2002), viii+1–123. 15, 16

[Bou06] ———, *Contact homology and homotopy groups of the space of contact structures*, Math. Res. Lett. **13** (2006), no. 1, 71–85. MR 2200047 (2007c:53122) 12

[Bou09] ———, *A survey of contact homology*, New perspectives and challenges in symplectic field theory, CRM Proc. Lecture Notes, vol. 49, Amer. Math. Soc., Providence, RI, 2009, pp. 45–71. MR 2555933 (2011a:53175) 13

[Bou11] ———, *Lectures on symplectic and contact homology*, Proceedings of the Gökova Geometry-Topology Conference 2010, Int. Press, Somerville, MA, 2011, pp. 1–20. MR 2931878

[BvK10] Frédéric Bourgeois and Otto van Koert, *Contact homology of left-handed stabilizations and plumbing of open books*, Commun. Contemp. Math. **12** (2010), no. 2, 223–263. MR 2646902 (2012d:53278) 13

[CDvK14] River Chiang, Fan Ding, and Otto van Koert, *Non-fillable invariant contact structures on principal circle bundles and left-handed twists*, Arxiv Preprint **arXiv:1404.1705** (2014), 1–13. 14

[CE12] Kai Cieliebak and Yakov Eliashberg, *From Stein to Weinstein and back*, American Mathematical Society Colloquium Publications, vol. 59, American Mathematical Society, Providence, RI, 2012, Symplectic geometry of affine complex manifolds. MR 3012475 12

[CGH14] Daniel Cristofaro-Gardiner and Michael Hutchings, *From one Reeb orbit to two*, Arxiv Preprint **arXiv:1202.4839v4** (2014), 1–13. 14

[CGHR15] Daniel Cristofaro-Gardiner, Michael Hutchings, and Vinicius Gripp Barros Ramos, *The asymptotics of ECH capacities*, Invent. Math. **199** (2015), no. 1, 187–214. MR 3294959 14

[CH13] Vincent Colin and Ko Honda, *Reeb vector fields and open book decompositions*, J. Eur. Math. Soc. (JEMS) **15** (2013), no. 2, 443–507. MR 3017043 15

[CL09] Kai Cieliebak and Janko Latschev, *The role of string topology in symplectic field theory*, New perspectives and challenges in symplectic field theory, CRM Proc. Lecture Notes, vol. 49, Amer. Math. Soc., Providence, RI, 2009, pp. 113–146. MR 2555935 (2010j:53187) 15

[CMP15] Roger Casals, Emmy Murphy, and Francisco Presas, *Geometric criteria for overtwistedness*, Arxiv Preprint **arXiv:1503.06221** (2015), 1–23. 13

[DGZ14] Max Dörner, Hansjörg Geiges, and Kai Zehmisch, *Open books and the Weinstein conjecture*, Q. J. Math. **65** (2014), no. 3, 869–885. MR 3261972 14

[Dra04] Dragomir L. Dragnev, *Fredholm theory and transversality for noncompact pseudoholomorphic maps in symplectizations*, Comm. Pure Appl. Math. **57** (2004), no. 6, 726–763. MR 2038115 (2004m:53153)

[EES05] Tobias Ekholm, John Etnyre, and Michael Sullivan, *The contact homology of Legendrian submanifolds in* $\mathbb{R}^{2n+1}$, J. Differential Geom. **71** (2005), no. 2, 177–305. MR 2197142 (2007f:53115) 1

[EGH00] Y. Eliashberg, A. Givental, and H. Hofer, *Introduction to symplectic field theory*, Geom. Funct. Anal. (2000), no. Special Volume, Part II, 560–673, GAFA 2000 (Tel Aviv, 1999). MR 1826267 (2002e:53136) 1, 2, 7, 11, 15, 16, 27, 29

[EH89] I. Ekeland and H. Hofer, *Symplectic topology and Hamiltonian dynamics*, Math. Z. **200** (1989), no. 3, 355–378. MR 978597 (90a:58046) 15

[EH90] Ivar Ekeland and Helmut Hofer, *Symplectic topology and Hamiltonian dynamics. II*, Math. Z. **203** (1990), no. 4, 553–567. MR 1044064 (91e:58053) 15

[EHK12] Tobias Ekholm, Ko Honda, and Tamás Kálmán, *Legendrian knots and exact Lagrangian cobordisms*, Arxiv Preprint **arXiv:1212.1519v3** (2012), 1–57. 7

[Ekh15] Tobias Ekholm, *Non-loose Legendrian spheres with trivial Contact Homology DGA*, Arxiv Preprint **arXiv:1502.04526** (2015), 1–21. 13

[EKP06] Yakov Eliashberg, Sang Seon Kim, and Leonid Polterovich, *Geometry of contact transformations and domains: orderability versus squeezing*, Geom. Topol. **10** (2006), 1635–1747. MR 2284048 (2007j:53111) 15

[Eli89] Y. Eliashberg, *Classification of overtwisted contact structures on 3-manifolds*, Invent. Math. **98** (1989), no. 3, 623–637. MR 1022310 (90k:53064) 13

[Eli90] Yakov Eliashberg, *Filling by holomorphic discs and its applications*, Geometry of low-dimensional manifolds, 2 (Durham, 1989), London Math. Soc. Lecture Note Ser., vol. 151, Cambridge Univ. Press, Cambridge, 1990, pp. 45–67. MR 1171908 (93g:53060)

[Eli91] ———, *On symplectic manifolds with some contact properties*, J. Differential Geom. **33** (1991), no. 1, 233–238. MR 1085141 (92g:57036)

[Eli98] ———, *Invariants in contact topology*, Proceedings of the International Congress of Mathematicians, Vol. II (Berlin, 1998), no. Extra Vol. II, 1998, pp. 327–338. MR 1648083 (2000a:57068) 1, 2, 11, 13

[Eli07] ———, *Symplectic field theory and its applications*, International Congress of Mathematicians. Vol. I, Eur. Math. Soc., Zürich, 2007, pp. 217–246. MR 2334192 (2008f:53119) 2

[EO15] Tobias Ekholm and Alexandru Oancea, *Symplectic and contact differential graded algebras*, Arxiv Preprint **arXiv:1506.01078v1** (2015), 1–58. 7

[Fab10] Oliver Fabert, *Contact homology of Hamiltonian mapping tori*, Comment. Math. Helv. **85** (2010), no. 1, 203–241. MR 2563686 (2011b:53215)

[FH93] A. Floer and H. Hofer, *Coherent orientations for periodic orbit problems in symplectic geometry*, Math. Z. **212** (1993), no. 1, 13–38. MR 1200162 (94m:58036) 27, 29

[FO99] Kenji Fukaya and Kaoru Ono, *Arnold conjecture and Gromov-Witten invariant*, Topology **38** (1999), no. 5, 933–1048. MR 1688434 (2000j:53116) 1

[FOOO09a] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono, *Lagrangian intersection Floer theory: anomaly and obstruction. Part I*, AMS/IP Studies in Advanced Mathematics, vol. 46, American Mathematical Society, Providence, RI, 2009. MR 2553465 (2011c:53217) 1

[FOOO09b] ______, *Lagrangian intersection Floer theory: anomaly and obstruction. Part II*, AMS/IP Studies in Advanced Mathematics, vol. 46, American Mathematical Society, Providence, RI, 2009. MR 2548482 (2011c:53218) 1

[FOOO12] Kenji Fukaya, Yong-Geun Oh, Hiroshi Ohta, and Kaoru Ono, *Technical details on Kuranishi structure and virtual fundamental chain*, Arxiv Preprint **arXiv:1209.4410v1** (2012), 1–257. 1

[FOOO15] ______, *Kuranishi structure, Pseudo-holomorphic curve, and Virtual fundamental chain: Part 1*, Arxiv Preprint **arXiv:1503.07631v1** (2015), 1–203. 1

[Gei97] Hansjörg Geiges, *Applications of contact surgery*, Topology **36** (1997), no. 6, 1193–1220. MR 1452848 (98d:57044)

[Ghi06] Paolo Ghiggini, *Infinitely many universally tight contact manifolds with trivial Ozsváth-Szabó contact invariants*, Geom. Topol. **10** (2006), 335–357 (electronic). MR 2224460 (2007a:57030) 14

[GK10] Viktor L. Ginzburg and Ely Kerman, *Homological resonances for Hamiltonian diffeomorphisms and Reeb flows*, Int. Math. Res. Not. IMRN (2010), no. 1, 53–68. MR 2576283 (2011b:53216) 14

[GM15] W. D. Gillam and Samouil Molcho, *Log differentiable spaces and manifolds with corners*, Arxiv Preprint **arXiv:1507.06752v1** (2015), 1–128. 30

[Gro85] M. Gromov, *Pseudoholomorphic curves in symplectic manifolds*, Invent. Math. **82** (1985), no. 2, 307–347. MR 809718 (87j:53053) 1, 2, 74

[Gut15] Jean Gutt, *The positive equivariant symplectic homology as an invariant for some contact manifolds*, Arxiv Preprint **arXiv:1503.01443** (2015), 1–34. 15

[GZ12] Hansjörg Geiges and Kai Zehmisch, *Symplectic cobordisms and the strong Weinstein conjecture*, Math. Proc. Cambridge Philos. Soc. **153** (2012), no. 2, 261–279. MR 2981927 14

[GZ14] ______, *The Weinstein conjecture for connected sums*, Arxiv Preprint **arXiv:1407.4307** (2014), 1–13. 14

[HN15] Michael Hutchings and Jo Nelson, *Cylindrical contact homology for dynamically convex contact forms in three dimensions*, Arxiv Preprint **arXiv:1407.2898** (2015), 1–27. 1

[Hof93] H. Hofer, *Pseudoholomorphic curves in symplectizations with applications to the Weinstein conjecture in dimension three*, Invent. Math. **114** (1993), no. 3, 515–563. MR 1244912 (94j:58064) 2, 14

[HT07] Michael Hutchings and Clifford Henry Taubes, *Gluing pseudoholomorphic curves along branched covered cylinders. I*, J. Symplectic Geom. **5** (2007), no. 1, 43–137. MR 2371184 (2008k:53201) 1, 14, 58

[HT09a] ———, *Gluing pseudoholomorphic curves along branched covered cylinders. II*, J. Symplectic Geom. **7** (2009), no. 1, 29–133. MR 2491716 (2010c:53130) 1, 14, 58

[HT09b] ———, *The Weinstein conjecture for stable Hamiltonian structures*, Geom. Topol. **13** (2009), no. 2, 901–941. MR 2470966 (2010h:53139) 14

[Hut02] Michael Hutchings, *An index inequality for embedded pseudoholomorphic curves in symplectizations*, J. Eur. Math. Soc. (JEMS) **4** (2002), no. 4, 313–361. MR 1941088 (2004b:53148) 1, 14

[Hut09] ———, *The embedded contact homology index revisited*, New perspectives and challenges in symplectic field theory, CRM Proc. Lecture Notes, vol. 49, Amer. Math. Soc., Providence, RI, 2009, pp. 263–297. MR 2555941 (2010k:53152) 1, 14

[Hut11] ———, *Quantitative embedded contact homology*, J. Differential Geom. **88** (2011), no. 2, 231–266. MR 2838266 15

[Hut15] ———, *Beyond ECH capacities*, Arxiv Preprint **arXiv:1409.1352v4** (2015), 1–42. 15

[HWZ95] H. Hofer, K. Wysocki, and E. Zehnder, *Properties of pseudo-holomorphic curves in symplectisations. II. Embedding controls and algebraic invariants*, Geom. Funct. Anal. **5** (1995), no. 2, 270–328. MR 1334869 (97e:58030) 2, 39

[HWZ96] ———, *Properties of pseudoholomorphic curves in symplectisations. I. Asymptotics*, Ann. Inst. H. Poincaré Anal. Non Linéaire **13** (1996), no. 3, 337–379. MR 1395676 (97e:58029) 2

[HWZ98a] ———, *Correction to: "Properties of pseudoholomorphic curves in symplectisations. I. Asymptotics" [Ann. Inst. H. Poincaré Anal. Non Linéaire* **13** *(1996), no. 3, 337–379; MR1395676 (97e:58029)]*, Ann. Inst. H. Poincaré Anal. Non Linéaire **15** (1998), no. 4, 535–538. MR 1632925 (99e:58026) 2

[HWZ98b] ———, *The dynamics on three-dimensional strictly convex energy surfaces*, Ann. of Math. (2) **148** (1998), no. 1, 197–289. MR 1652928 (99m:58089)

[HWZ99] ———, *Properties of pseudoholomorphic curves in symplectizations. III. Fredholm theory*, Topics in nonlinear analysis, Progr. Nonlinear Differential Equations Appl., vol. 35, Birkhäuser, Basel, 1999, pp. 381–475. MR 1725579 (2001e:53092) 2

[HWZ02] ———, *Finite energy cylinders of small area*, Ergodic Theory Dynam. Systems **22** (2002), no. 5, 1451–1486. MR 1934144 (2003k:53121) 2, 72, 74

[HWZ03] ———, *Finite energy foliations of tight three-spheres and Hamiltonian dynamics*, Ann. of Math. (2) **157** (2003), no. 1, 125–255. MR 1954266 (2004a:53108) 14

[HWZ07] H. Hofer, K. Wysocki, and E. Zehnder, *A general Fredholm theory. I. A splicing-based differential geometry*, J. Eur. Math. Soc. (JEMS) **9** (2007), no. 4, 841–876. MR 2341834 (2008m:53202) 1

[HWZ09a] Helmut Hofer, Krzysztof Wysocki, and Eduard Zehnder, *A general Fredholm theory. II. Implicit function theorems*, Geom. Funct. Anal. **19** (2009), no. 1, 206–293. MR 2507223 (2010g:53174) 1

[HWZ09b] Helmut Hofer, Kris Wysocki, and Eduard Zehnder, *A general Fredholm theory. III. Fredholm functors and polyfolds*, Geom. Topol. **13** (2009), no. 4, 2279–2387. MR 2515707 (2010h:53138) 1

[HWZ10a] H. Hofer, K. Wysocki, and E. Zehnder, *Integration theory on the zero sets of polyfold Fredholm sections*, Math. Ann. **346** (2010), no. 1, 139–198. MR 2558891 (2011c:53222) 1

[HWZ10b] Helmut Hofer, Kris Wysocki, and Eduard Zehnder, *sc-smoothness, retractions and new models for smooth spaces*, Discrete Contin. Dyn. Syst. **28** (2010), no. 2, 665–788. MR 2644764 (2011k:58006) 1

[HWZ11] Helmut Hofer, Kris Wysocki, and Eduard Zehnder, *Applications of polyfold theory I: The polyfolds of Gromov–Witten theory*, Arxiv Preprint **arXiv:1107.2097** (2011), 1–205. 1

[HWZ14] Helmut Hofer, Kris Wysocki, and Eduard Zehnder, *Polyfold and Fredholm Theory I: Basic Theory in M-polyfolds*, Arxiv Preprint **arXiv:1407.3185v1** (2014), 1–246. 1

[Joy12] Dominic Joyce, *D-manifolds and d-orbifolds: a theory of derived differential geometry*, Book Manuscript (2012), 1–768. 1

[Joy14] Dominic Joyce, *An introduction to d-manifolds and derived differential geometry*, Moduli spaces, London Math. Soc. Lecture Note Ser., vol. 411, Cambridge Univ. Press, Cambridge, 2014, pp. 230–281. MR 3221297 1

[Joy15a] ———, *A generalization of manifolds with corners*, Arxiv Preprint **arXiv:1501.00401v1** (2015), 1–95. 30

[Joy15b] Dominic Joyce, *A new definition of Kuranishi space*, Arxiv Preprint **arXiv:1409.6908v2** (2015), 1–192. 1

[KM15] Chris Kottke and Richard B. Melrose, *Generalized blow-up of corners and fiber products*, Trans. Amer. Math. Soc. **367** (2015), no. 1, 651–705. MR 3271273 30

[LM85] Robert B. Lockhart and Robert C. McOwen, *Elliptic differential operators on noncompact manifolds*, Ann. Scuola Norm. Sup. Pisa Cl. Sci. (4) **12** (1985), no. 3, 409–447. MR 837256 (87k:58266) 26

[LT98a] Jun Li and Gang Tian, *Virtual moduli cycles and Gromov-Witten invariants of general symplectic manifolds*, Topics in symplectic 4-manifolds (Irvine, CA, 1996), First Int. Press Lect. Ser., I, Int. Press, Cambridge, MA, 1998, pp. 47–83. MR 1635695 (2000d:53137) 1

[LT98b] Gang Liu and Gang Tian, *Floer homology and Arnold conjecture*, J. Differential Geom. **49** (1998), no. 1, 1–74. MR 1642105 (99m:58047) 1

[LW11] Janko Latschev and Chris Wendl, *Algebraic torsion in contact manifolds*, Geom. Funct. Anal. **21** (2011), no. 5, 1144–1195, With an appendix by Michael Hutchings. MR 2846386 14

[MNW13] Patrick Massot, Klaus Niederkrüger, and Chris Wendl, *Weak and strong fillability of higher dimensional contact manifolds*, Invent. Math. **192** (2013), no. 2, 287–373. MR 3044125 13, 14

[MS04] Dusa McDuff and Dietmar Salamon, *J-holomorphic curves and symplectic topology*, American Mathematical Society Colloquium Publications, vol. 52, American Mathematical Society, Providence, RI, 2004. MR 2045629 (2004m:53154) 39, 74

[Mul94] Marie-Paule Muller, *Gromov's Schwarz lemma as an estimate of the gradient for holomorphic curves*, Holomorphic curves in symplectic geometry, Progr. Math., vol. 117, Birkhäuser, Basel, 1994, pp. 217–231. MR 1274931 74

[MW15a] Dusa McDuff and Katrin Wehrheim, *The fundamental class of smooth Kuranishi atlases with trivial isotropy*, Arxiv Preprint **arXiv:1508.01560v2** (2015), 1–151. 1

[MW15b] ———, *Smooth Kuranishi atlases with isotropy*, Arxiv Preprint **arXiv:1508.01556v2** (2015), 1–69. 1

[MW15c] ———, *The topology of Kuranishi atlases*, Arxiv Preprint **arXiv:1508.01844v2** (2015), 1–71. 1

[Nie06] Klaus Niederkrüger, *The plastikstufe—a generalization of the overtwisted disk to higher dimensions*, Algebr. Geom. Topol. **6** (2006), 2473–2508. MR 2286033 (2007k:57053) 13

[NR11] Klaus Niederkrüger and Ana Rechtman, *The Weinstein conjecture in the presence of submanifolds having a Legendrian foliation*, J. Topol. Anal. **3** (2011), no. 4, 405–421. MR 2887669 14

[NvK07] Klaus Niederkrüger and Otto van Koert, *Every contact manifolds can be given a nonfillable contact structure*, Int. Math. Res. Not. IMRN (2007), no. 23, Art. ID rnm115, 22. MR 2380008 (2008j:53145)

[NW11] Klaus Niederkrüger and Chris Wendl, *Weak symplectic fillings and holomorphic curves*, Ann. Sci. Éc. Norm. Supér. (4) **44** (2011), no. 5, 801–853. MR 2931519 14

[OS05] Peter Ozsváth and Zoltán Szabó, *Heegaard Floer homology and contact structures*, Duke Math. J. **129** (2005), no. 1, 39–61. MR 2153455 (2006b:57043) 13

[Par15] John Pardon, *An algebraic approach to virtual fundamental cycles on moduli spaces of pseudo-holomorphic curves*, Geom. Topol. (to appear) **arXiv:1309.2370v5** (2015), 1–205. 1, 2, 16, 30, 33, 41, 42, 43, 44, 50, 51, 52, 56, 57, 58, 69, 74

[Rua99] Yongbin Ruan, *Virtual neighborhoods and pseudo-holomorphic curves*, Proceedings of 6th Gökova Geometry-Topology Conference, vol. 23, 1999, pp. 161–231. MR 1701645 (2002b:53138) 1

[Tau07] Clifford Henry Taubes, *The Seiberg-Witten equations and the Weinstein conjecture*, Geom. Topol. **11** (2007), 2117–2202. MR 2350473 (2009b:57055) 14

[Ust99] Ilya Ustilovsky, *Infinitely many contact structures on S^{4m+1}*, Internat. Math. Res. Notices (1999), no. 14, 781–791. MR 1704176 (2000f:57028) 15

[Vau15] Anne Vaugon, *On growth rate and contact homology*, Algebr. Geom. Topol. **15** (2015), no. 2, 623–666. MR 3342671 14

[vK08] Otto van Koert, *Contact homology of Brieskorn manifolds*, Forum Math. **20** (2008), no. 2, 317–339. MR 2395966 (2009d:53127) 15

[Wan14] Andy Wand, *Tightness is preserved by Legendrian surgery*, Arxiv Preprint **arXiv:1404.1705** (2014), 1–13. 13

[Wei79] Alan Weinstein, *On the hypotheses of Rabinowitz' periodic orbit theorems*, J. Differential Equations **33** (1979), no. 3, 353–358. MR 543704 (81a:58030b) 14

[Yau04] Mei-Lin Yau, *Cylindrical contact homology of subcritical Stein-fillable contact manifolds*, Geom. Topol. **8** (2004), 1243–1280. MR 2087083 (2005g:53172) 15

[Yau06] ———, *Vanishing of the contact homology of overtwisted contact 3-manifolds*, Bull. Inst. Math. Acad. Sin. (N.S.) **1** (2006), no. 2, 211–229, With an appendix by Yakov Eliashberg. MR 2230587 (2007a:57032) 13

[Yau09] ———, *Cylindrical contact homology of a Dehn twist*, Internat. J. Math. **20** (2009), no. 12, 1479–1525. MR 2597305 (2012d:53279) 15